HP ProLiant SB460c SAN Gateway Storage Server

user guide

Part number: AN583-96001
First edition: September 2008

Legal and notice information

Contents

Figures

Tables

About this guide

This guide provides information about configuring, managing, and troubleshooting the HP ProLiant SB460c SAN Gateway Storage Server.

Intended audience

This guide is intended for technical professionals with knowledge of:

- Microsoft® administrative procedures
- System and storage configurations

Related documentation

The following documents [and websites] provide related information:

- *HP Integrated Lights-Out 2 User Guide*
- *HP ProLiant Lights-Out 100 Remote Management User Guide*

You can find these documents from the Manuals page of the HP Business Support Center website:

http://www.hp.com/support/manuals

In the Storage section, click **Disk Storage Systems** and then select your product.

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Blue text: Table 1	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values

Convention	Element
Monospace, italic text	• Code variables • Command variables
Monospace, bold text	Emphasized monospace text

⚠ WARNING!

Indicates that failure to follow directions could result in bodily harm or death.

⚠ CAUTION:

Indicates that failure to follow directions could result in damage to equipment or data.

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

Rack stability

Rack stability protects personnel and equipment.

⚠ WARNING!

To reduce the risk of personal injury or damage to equipment:

- Extend leveling jacks to the floor.
- Ensure that the full weight of the rack rests on the leveling jacks.
- Install stabilizing feet on the rack.
- In multiple-rack installations, fasten racks together securely.
- Extend only one rack component at a time. Racks can become unstable if more than one component is extended.

HP technical support

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Customer self repair

HP customer self repair (CSR) programs allow you to repair your StorageWorks product. If a CSR part needs replacing, HP ships the part directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HP-authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider. For North America, see the CSR website:

http://www.hp.com/go/selfrepair

Product warranties

For information about HP StorageWorks product warranties, see the warranty information website:

http://www.hp.com/go/storagewarranty

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/e-updates

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

HP websites

For additional information, see the following HP websites:

- http://www.hp.com
- http://www.hp.com/go/storage
- http://www.hp.com/service_locator
- http://www.hp.com/support/manuals
- http://www.hp.com/support/downloads

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to storagedocsFeedback@hp.com. All submissions become the property of HP.

1 Storage management overview

This chapter provides an overview of some of the components that make up the storage structure of the HP ProLiant Storage Server.

Storage management elements

Storage is divided into four major divisions:

- Physical storage elements
- Logical storage elements
- File system elements
- File sharing elements

Each of these elements is composed of the previous level's elements.

Storage management example

Figure 1 depicts many of the storage elements that one would find on a storage device. The following sections provide an overview of the storage elements.

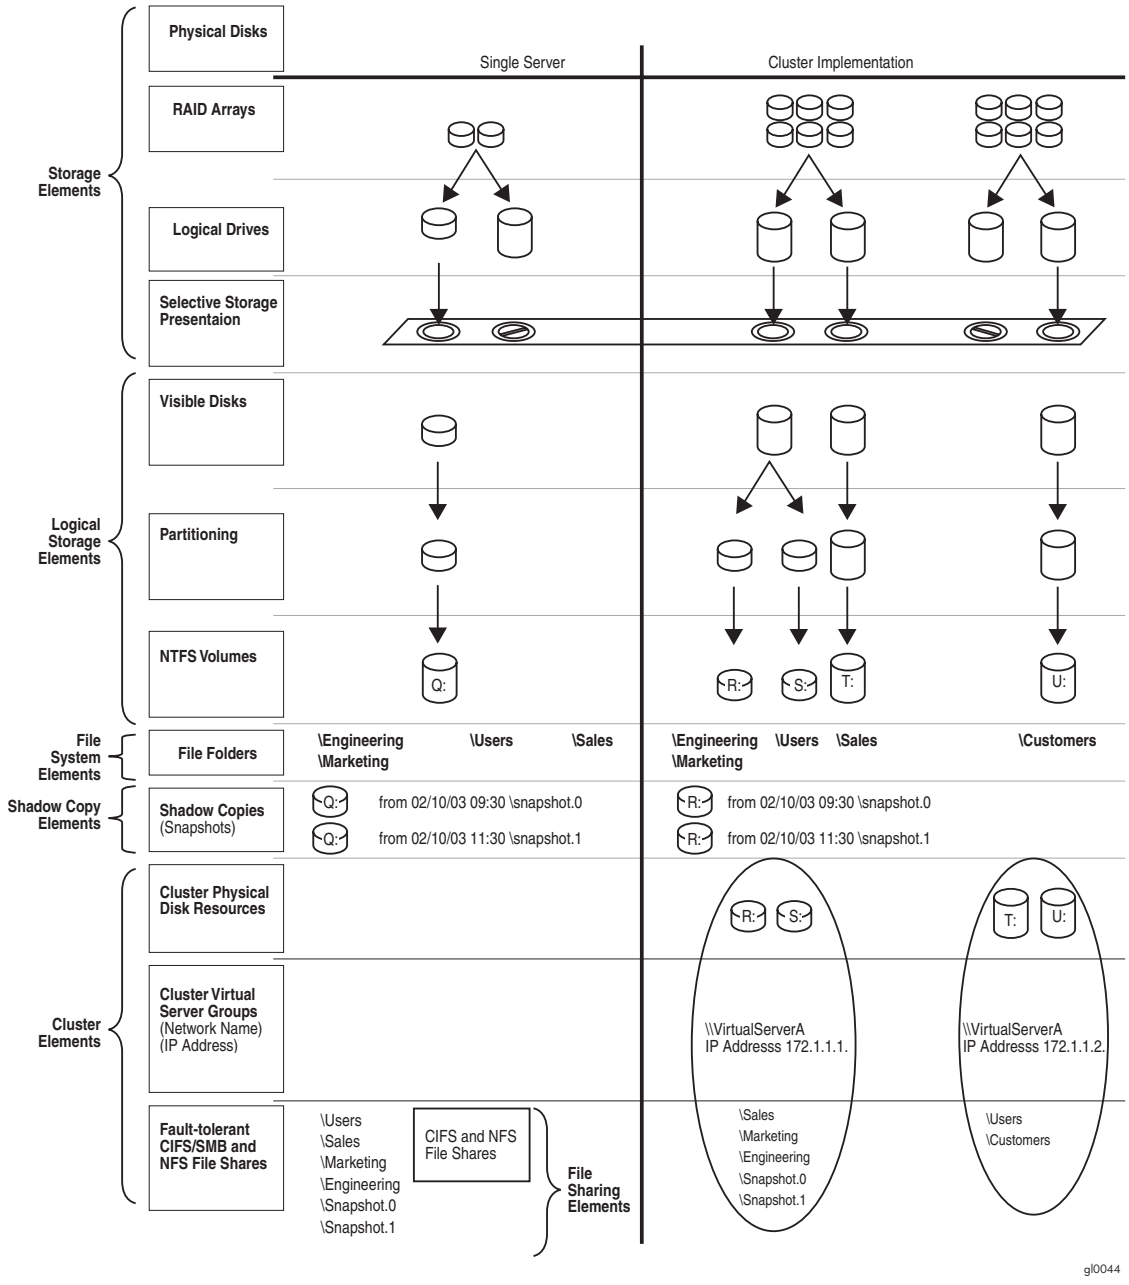

Figure 1 Storage management process example

Physical storage elements

The lowest level of storage management occurs at the physical drive level. Minimally, choosing the best disk carving strategy includes the following policies:

- Analyze current corporate and departmental structure.
- Analyze the current file server structure and environment.
- Plan properly to ensure the best configuration and use of storage.
  - Determine the desired priority of fault tolerance, performance, and storage capacity.
  - Use the determined priority of system characteristics to determine the optimal striping policy and RAID level.

- Include the appropriate number of physical drives in the arrays to create logical storage elements of desired sizes.

Arrays

See Figure 2. With an array controller installed in the system, the capacity of several physical drives (P1–P3) can be logically combined into one or more logical units (L1) called arrays. When this is done, the read/write heads of all the constituent physical drives are active simultaneously, dramatically reducing the overall time required for data transfer.

NOTE:

Depending on the storage server model, array configuration may not be possible or necessary.

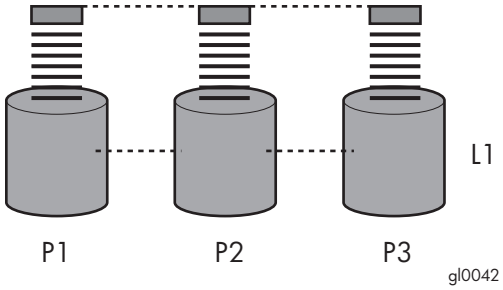

Figure 2 Configuring arrays from physical drives

Because the read/write heads are simultaneously active, the same amount of data is written to each drive during any given time interval. Each unit of data is termed a block. The blocks form a set of data stripes over all the hard drives in an array, as shown in Figure 3.

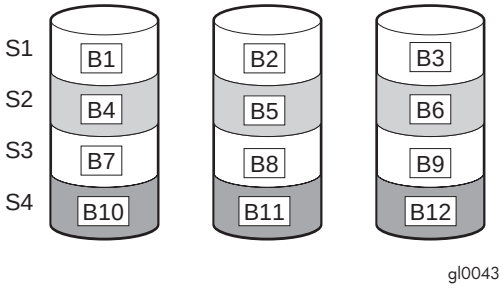

Figure 3 RAID 0 (data striping) (S1-S4) of data blocks (B1-B12)

For data in the array to be readable, the data block sequence within each stripe must be the same. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each hard drive in a given array contains the same number of data blocks.

NOTE:

If one hard drive has a larger capacity than other hard drives in the same array, the extra capacity is wasted because it cannot be used by the array.

Fault tolerance

Drive failure, although rare, is potentially catastrophic. For example, using simple striping as shown in Figure 3, failure of any hard drive leads to failure of all logical drives in the same array, and hence to data loss.

To protect against data loss from hard drive failure, storage servers should be configured with fault tolerance. HP recommends adhering to RAID 5 configurations.

The table below summarizes the important features of the different kinds of RAID supported by the Smart Array controllers. The decision chart in the following table can help determine which option is best for different situations.

Table 2 Summary of RAID methods

	RAID 0 Striping (no fault tolerance)	RAID 1+0 Mirroring	RAID 5 Distributed Data Guarding	RAID 6 (ADG)
Maximum number of hard drives	N/A	N/A	14	Storage system dependent
Tolerant of single hard drive failure?	No	Yes	Yes	Yes
Tolerant of multiple simultaneous hard drive failures?	No	If the failed drives are not mirrored to each other	No	Yes (two drives can fail)

Online spares

Further protection against data loss can be achieved by assigning an online spare (or hot spare) to any configuration except RAID 0. This hard drive contains no data and is contained within the same storage subsystem as the other drives in the array. When a hard drive in the array fails, the controller can then automatically rebuild information that was originally on the failed drive onto the online spare. This quickly restores the system to full RAID level fault tolerance protection. However, unless RAID Advanced Data Guarding (ADG) is being used, which can support two drive failures in an array, in the unlikely event that a third drive in the array should fail while data is being rewritten to the spare, the logical drive still fails.

Logical storage elements

Logical storage elements consist of those components that translate the physical storage elements to file system elements. The storage server uses the Window Disk Management utility to manage the various types of disks presented to the file system. There are two types of LUN presentation: basic disk and dynamic disk. Each of these types of disk has special features that enable different types of management.

Logical drives (LUNs)

While an array is a physical grouping of hard drives, a logical drive consists of components that translate physical storage elements into file system elements.

It is important to note that a LUN may span all physical drives within a storage controller subsystem, but cannot span multiple storage controller subsystems.

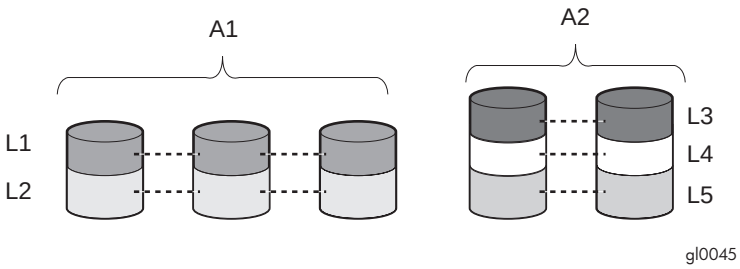

Figure 4 Two arrays (A1, A2) and five logical drives (L1 through L5) spread over five physical drives

NOTE:

This type of configuration may not apply to all storage servers and serves only as an example.

Through the use of basic disks, you can create primary partitions or extended partitions. Partitions can only encompass one LUN. Through the use of dynamic disks, you can create volumes that span multiple LUNs. You can use the Windows Disk Management utility to convert disks to dynamic and back to basic and to manage the volumes residing on dynamic disks. Other options include the ability to delete, extend, mirror, and repair these elements.

Partitions

Partitions exist as either primary partitions or extended partitions and can be composed of only one basic disk no larger than 2 TB. Basic disks can also only contain up to four primary partitions, or three primary partitions and one extended partition. In addition, the partitions on them cannot be extended beyond the limits of a single LUN. Extended partitions allow the user to create multiple logical drives. These partitions or logical disks can be assigned drive letters or be used as mount points on existing disks. If mount points are used, it should be noted that Services for UNIX (SFU) does not support mount points at this time. The use of mount points in conjunction with NFS shares is not supported.

Volumes

When planning dynamic disks and volumes, there is a limit to the amount of growth a single volume can undergo. Volumes are limited in size and can have no more than 32 separate LUNs, with each LUN not exceeding 2 terabytes (TB), and volumes totaling no more than 64 TB of disk space.

The RAID level of the LUNs included in a volume must be considered. All of the units that make up a volume should have the same high-availability characteristics. In other words, the units should all be of the same RAID level. For example, it would not be a good practice to include both a RAID 1+0 and a RAID 5 array in the same volume set. By keeping all the units the same, the entire volume retains the same performance and high-availability characteristics, making managing and maintaining the volume much easier. If a dynamic disk goes offline, the entire volume dependent on the one or more dynamic disks is unavailable. There could be a potential for data loss depending on the nature of the failed LUN.

Volumes are created out of the dynamic disks, and can be expanded on the fly to extend over multiple dynamic disks if they are spanned volumes. However, after a type of volume is selected, it cannot be altered. For example, a spanning volume cannot be altered to a mirrored volume without deleting and recreating the volume, unless it is a simple volume. Simple volumes can be mirrored or converted to spanned volumes. Fault-tolerant disks cannot be extended. Therefore, selection of the volume type

is important. The same performance characteristics on numbers of reads and writes apply when using fault-tolerant configurations, as is the case with controller-based RAID. These volumes can also be assigned drive letters or be mounted as mount points off existing drive letters.

The administrator should carefully consider how the volumes will be carved up and what groups or applications will be using them. For example, putting several storage-intensive applications or groups into the same dynamic disk set would not be efficient. These applications or groups would be better served by being divided up into separate dynamic disks, which could then grow as their space requirements increased, within the allowable growth limits.

NOTE:

Dynamic disks cannot be used for clustering configurations because Microsoft Cluster only supports basic disks.

File system elements

File system elements are composed of the folders and subfolders that are created under each logical storage element (partitions, logical disks, and volumes). Folders are used to further subdivide the available file system, providing another level of granularity for management of the information space. Each of these folders can contain separate permissions and share names that can be used for network access. Folders can be created for individual users, groups, projects, and so on.

File sharing elements

The storage server supports several file sharing protocols, including Distributed File System (DFS), Network File System (NFS), File Transfer Protocol (FTP), Hypertext Transfer Protocol (HTTP), and Microsoft Server Message Block (SMB). On each folder or logical storage element, different file sharing protocols can be enabled using specific network names for access across a network to a variety of clients. Permissions can then be granted to those shares based on users or groups of users in each of the file sharing protocols.

Volume Shadow Copy Service overview

The Volume Shadow Copy Service (VSS) provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. VSS supports 64 shadow copies per volume.

Shadow Copies of Shared Folders resides within this infrastructure, and helps alleviate data loss by creating shadow copies of files or folders that are stored on network file shares at pre-determined time intervals. In essence, a shadow copy is a previous version of the file or folder at a specific point in time.

By using shadow copies, a storage server can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer.

Shadow copies should not replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. For example, shadow copies cannot protect against data loss due to media failures; however, recovering data from shadow copies can reduce the number of times needed to restore data from tape.

Using storage elements

The last step in creating the element is determining its drive letter or mount point and formatting the element. Each element created can exist as a drive letter, assuming one is available, and/or as mount points on an existing folder or drive letter. Either method is supported. However, mount points cannot be used for shares that will be shared using Microsoft Services for Unix. They can be set up with both but the use of the mount point in conjunction with NFS shares causes instability with the NFS shares.

Formats consist of NTFS, FAT32, and FAT. All three types can be used on the storage server. However, VSS can only use volumes that are NTFS formatted. Also, quota management is possible only on NTFS.

Clustered server elements

Select storage servers support clustering. The HP ProLiant Storage Server supports several file sharing protocols including DFS, NFS, FTP, HTTP, and Microsoft SMB. Only NFS, FTP, and Microsoft SMB are cluster-aware protocols. HTTP can be installed on each node but the protocols cannot be set up through cluster administrator, and they will not fail over during a node failure.

△ CAUTION:

AppleTalk shares should not be created on clustered resources as this is not supported by Microsoft Clustering, and data loss may occur.

Network names and IP address resources for the clustered file share resource can also be established for access across a network to a variety of clients. Permissions can then be granted to those shares based on users or groups of users in each of the file sharing protocols.

Network adapter teaming

Network adapter teaming is software-based technology used to increase a server's network availability and performance. Teaming enables the logical grouping of physical adapters in the same server (regardless of whether they are embedded devices or Peripheral Component Interconnect (PCI) adapters) into a virtual adapter. This virtual adapter is seen by the network and server-resident network-aware applications as a single network connection.

Management tools

HP Systems Insight Manager

HP SIM is a web-based application that allows system administrators to accomplish normal administrative tasks from any remote location, using a web browser. HP SIM provides device management capabilities that consolidate and integrate management data from HP and third-party devices.

ⓘ IMPORTANT:

You must install and use HP SIM to benefit from the Pre-Failure Warranty for processors, SAS and SCSI hard drives, and memory modules.

For additional information, refer to the Management CD in the HP ProLiant Essentials Foundation Pack or the HP SIM website (http://www.hp.com/go/hpsim).

Management Agents

Management Agents provide the information to enable fault, performance, and configuration management. The agents allow easy manageability of the server through HP SIM software, and thirdparty SNMP management platforms. Management Agents are installed with every SmartStart assisted installation or can be installed through the HP PSP. The Systems Management homepage provides status and direct access to in-depth subsystem information by accessing data reported through the Management Agents. For additional information, refer to the Management CD in the HP ProLiant Essentials Foundation Pack or the HP website (http://www.hp.com/servers/manage).

2 File server management

This chapter begins by identifying file services in Windows Storage Server 2003 R2. The remainder of the chapter describes the many tasks and utilities that play a role in file server management.

File services features in Windows Storage Server 2003 R2

Storage Manager for SANs

The Storage Manager for SANs (also called Simple SAN) snap-in enables you to create and manage the LUNs that are used to allocate space on storage arrays. Storage Manager for SANs can be used on SANs that support Virtual Disk Server (VDS). It can be used in both Fibre Channel and iSCSI environments.

For more information on Storage Manager for SANs, see the online help. A Microsoft document titled *Storage Management in Windows Storage Server 2003 R2: File Server Resource Manager and Storage Manager for Storage Area Networks* is available at http://download.microsoft.com/download/7/4/7/7472bf9b-3023-48b7-87be-d2cedc38f15a/WS03R2_Storage_Management.doc.

NOTE:

Storage Manager for SANs is only available on Standard and Enterprise editions of Windows Storage Server 2003 R2.

Single Instance Storage

Single Instance Storage (SIS) provides a copy-on-write link between multiple files. Disk space is recovered by reducing the amount of redundant data stored on a server. If a user has two files sharing disk storage by using SIS, and someone modifies one of the files, users of the other files do not see the changes. The underlying shared disk storage that backs SIS links is maintained by the system and is only deleted if all the SIS links pointing to it are deleted. SIS automatically determines that two or more files have the same content and links them together.

NOTE:

Single Instance Storage is only available on Standard and Enterprise editions of Windows Storage Server 2003 R2.

File Server Resource Manager

File Server Resource Manager is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. By using File Server Resource Manager,

administrators can place quotas on volumes, actively screen files and folders, and generate comprehensive storage reports.

By using File Server Resource Manager, you can perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and to generate notifications when the quota limits are approached and exceeded.
- Create file screens to screen the files that users can save on volumes and in folders and to send notifications when users attempt to save blocked files.
- Schedule periodic storage reports that allow users to identify trends in disk usage and to monitor attempts to save unauthorized files, or generate the reports on demand.

Windows SharePoint Services

Windows SharePoint Services is an integrated set of collaboration and communication services designed to connect people, information, processes, and systems, within and beyond the organization firewall.

NOTE:

Windows SharePoint Services is only available on Standard and Enterprise editions of Windows Storage Server 2003 R2.

HP Storage Server Management Console

The HP Storage Server Management Console is a user interface in Windows Storage Server 2003 R2 and Windows Unified Data Storage Server 2003 that provides one place to manage files or print serving components. The console is accessible using Remote Desktop or a web browser.

The Storage Management page provides a portal to:

- File Server Resource Manager
- DFS Management
- Disk and Volume Management
- Single Instance Storage
- Indexing Service
- MSNFS (under Share folder)
- Cluster Management (under “Utilities”)

The Share Folder Management page provides a portal to Shared Folders, consisting of:

- Shares
- Sessions
- Open files

File services management

Information about the storage server in a SAN environment is provided in the *HP ProLiant Storage Server SAN Connection and Management* document located on the HP web site at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00663737/c00663737.pdf.

Configurable and pre-configured storage

Certain storage servers ship with storage configured only for the operating system. The administrator must configure data storage for the storage server. Other storage servers ship with pre-configured storage for data. Depending on the type of storage server purchased, additional storage configuration is required.

Configuring additional storage involves creating arrays, logical disks, and volumes. Table 3 shows the general task areas to be performed as well as the utilities needed to configure storage for an HP Smart Array-based storage server.

Table 3 Tasks and utilities needed for storage server configuration

Task	Storage management utility
Create disk arrays	HP Array Configuration Utility or Storage Manager
Create logical disks from the array space	HP Array Configuration Utility or Storage Manager
Verify newly created logical disks	Windows Disk Management
Create a volume on the new logical disk	Windows Disk Management

- Create disk arrays—On storage servers with configurable storage, physical disks can be arranged as RAID arrays for fault tolerance and enhanced performance, and then segmented into logical disks of appropriate sizes for particular storage needs. These logical disks then become the volumes that appear as drives on the storage server.

△ CAUTION:

The first two logical drives are configured for the storage server operating system and should not be altered in any manner. If the first two logical drives are altered, the system recovery process may not function properly when using the System Recovery DVD. Do not tamper with the "DON'T ERASE" or local C: volume. These are reserved volumes and must be maintained as they exist.

The fault tolerance level depends on the amount of disks selected when the array was created. A minimum of two disks is required for RAID 0+1 configuration, three disks for a RAID 5 configuration, and four disks for a RAID 6 (ADG) configuration.

- Create logical disks from the array space—Select the desired fault tolerance, stripe size, and size of the logical disk.
- Verify newly created logical disks—Verify that disks matching the newly created sizes are displayed.
- Create a volume on the new logical disk—Select a drive letter and enter a volume label, volume size, allocation unit size, and mount point (if desired).

Storage management utilities

The storage management utilities preinstalled on the storage server include the HP Array Configuration Utility (ACU).

Array management utilities

Storage devices for RAID arrays and LUNs are created and managed using the array management utilities mentioned previously. For HP Smart Arrays use the ACU.

NOTE:

The ACU is used to configure and manage array-based storage. Software RAID-based storage servers use Microsoft Disk Manager to manage storage. You need administrator or root privileges to run the ACU.

Array Configuration Utility

The HP ACU supports the Smart Array controllers and hard drives installed on the storage server.

To open the ACU from the storage server desktop:

NOTE:

If this is the first time that the ACU is being run, you will be prompted to select the Execution Mode for ACU. Selecting Local Application Mode allows you to run the ACU from a Remote Desktop, remote console, or storage server web access mode. Remote service mode allows you to access the ACU from a remote browser.

1. Select **Start > Programs > HP Management Tools > Array Configuration Utility**.
2. If the Execution Mode for ACU is set to Remote Mode, log on to the HP System Management Homepage. The default user name is **administrator** and the default password is **hpinvent**.

To open the ACU in browser mode:

NOTE:

Confirm that the ACU Execution Mode is set to remote service.

1. Open a browser and enter the server name or IP address of the destination server. For example, `http://servername:2301` or `http://192.0.0.1:2301`.
2. Log on to the HP System Management Homepage. The default user name is **administrator** and the default password is **hpinvent**.
3. Click **Array Configuration Utility** on the left side of the window. The ACU opens and identifies the controllers that are connected to the system.

Some ACU guidelines to consider:

- Do not modify the first two logical drives of the storage server; they are configured for the storage server operating system.
- Spanning more than 14 disks with a RAID 5 volume is not recommended.
- Designate spares for RAID sets to provide greater protection against failures.
- RAID sets cannot span controllers.
- A single array can contain multiple logical drives of varying RAID settings.
- Extending and expanding arrays and logical drives is supported.

The *HP Array Configuration Utility User Guide* is available for download at http://www.hp.com/support/manuals.

Disk Management utility

The Disk Management tool is a system utility for managing hard disks and the volumes, or partitions, that they contain. Disk Management is used to initialize disks, create volumes, format volumes with the FAT, FAT32, or NTFS file systems, and create fault-tolerant disk systems. Most disk-related tasks can be performed in Disk Management without restarting the system or interrupting users. Most configuration changes take effect immediately. A complete online help facility is provided with the Disk Management utility for assistance in using the product.

NOTE:

- When the Disk Management utility is accessed through a Remote Desktop connection, this connection can only be used to manage disks and volumes on the server. Using the Remote Desktop connection for other operations during an open session closes the session.
- When closing Disk Management through a Remote Desktop connection, it may take a few moments for the remote session to log off.

Guidelines for managing disks and volumes

- The first two logical drives are configured for the storage server operating system and should not be altered in any manner. If the first two logical drives are altered, the system recovery process may not function properly when using the System Recovery DVD. Do not tamper with the "DON'T ERASE" or local C: volume. These are reserved volumes and must be maintained as they exist.
- HP does not recommend spanning array controllers with dynamic volumes. The use of software RAID-based dynamic volumes is not recommended. Use the array controller instead; it is more efficient.
- Use meaningful volume labels with the intended drive letter embedded in the volume label, if possible. (For example, volume e: might be named "Disk E:.") Volume labels often serve as the only means of identification.
- Record all volume labels and drive letters in case the system needs to be restored.
- When managing basic disks, only the last partition on the disk can be extended unless the disk is changed to dynamic.
- Basic disks can be converted to dynamic, but cannot be converted back to basic without deleting all data on the disk.
- Basic disks can contain up to four primary partitions (or three primary partitions and one extended partition).
- Format drives with a 16 K allocation size for best support of shadow copies, performance, and defragmentation.
- NTFS formatted drives are recommended because they provide the greatest level of support for shadow copies, encryption, and compression.
- Only basic disks can be formatted as FAT or FAT32.
- Read the online Disk Management help found in the utility.

Scheduling defragmentation

Defragmentation is the process of analyzing local volumes and consolidating fragmented files and folders so that each occupies a single, contiguous space on the volume. This improves file system performance. Because defragmentation consolidates files and folders, it also consolidates the free space on a volume. This reduces the likelihood that new files will be fragmented.

Defragmentation for a volume can be scheduled to occur automatically at convenient times. Defragmentation can also be done once, or on a recurring basis.

NOTE:

Scheduling defragmentation to run no later than a specific time prevents the defragmentation process from running later than that time. If the defragmentation process is running when the time is reached, the process is stopped. This setting is useful to ensure that the defragmentation process ends before the demand for server access is likely to increase.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger during the format. Otherwise defragmentation registers as a change by the Shadow Copy process. This increase in the number of changes forces Shadow Copy to delete snapshots as the limit for the cache file is reached.

CAUTION:

Allocation unit size cannot be altered without reformatting the drive. Data on a reformatted drive cannot be recovered.

For more information about disk defragmentation, read the online help.

Disk quotas

Disk quotas track and control disk space use in volumes.

NOTE:

To limit the size of a folder or share, see "Quota management" on page 53 .

Configure the volumes on the server to perform the following tasks:

- Prevent further disk space use and log an event when a user exceeds a specified disk space limit.
- Log an event when a user exceeds a specified disk space warning level.

When enabling disk quotas, it is possible to set both the disk quota limit and the disk quota warning level. The disk quota limit specifies the amount of disk space a user is allowed to use. The warning level specifies the point at which a user is nearing his or her quota limit. For example, a user's disk quota limit can be set to 50 megabytes (MB), and the disk quota warning level to 45 MB. In this case, the user can store no more than 50 MB on the volume. If the user stores more than 45 MB on the volume, the disk quota system logs a system event.

In addition, it is possible to specify that users can exceed their quota limit. Enabling quotas and not limiting disk space use is useful to still allow users access to a volume, but track disk space use on a per-user basis. It is also possible to specify whether or not to log an event when users exceed either their quota warning level or their quota limit.

When enabling disk quotas for a volume, volume usage is automatically tracked from that point forward, but existing volume users have no disk quotas applied to them. Apply disk quotas to existing volume users by adding new quota entries on the Quota Entries page.

NOTE:

When enabling disk quotas on a volume, any users with write access to the volume who have not exceeded their quota limit can store data on the volume. The first time a user writes data to a quota-enabled volume, default values for disk space limit and warning level are automatically assigned by the quota system.

For more information about disk quotas, read the online help.

Adding storage

Expansion is the process of adding physical disks to an array that has already been configured. Extension is the process of adding new storage space to an existing logical drive on the same array, usually after the array has been expanded.

Storage growth may occur in three forms:

- Extend unallocated space from the original logical disks or LUNs.
- Alter LUNs to contain additional storage.
- Add new LUNs to the system.

The additional space is then extended through a variety of means, depending on which type of disk structure is in use.

NOTE:

This section addresses only single storage server node configurations. If your server has Windows Storage Server 2003 R2 Enterprise Edition, see the Cluster Administration chapter for expanding and extending storage in a cluster environment.

Expanding storage

Expansion is the process of adding physical disks to an array that has already been configured. The logical drives (or volumes) that exist in the array before the expansion takes place are unchanged, because only the amount of free space in the array changes. The expansion process is entirely independent of the operating system.

NOTE:

See your storage array hardware user documentation for further details about expanding storage on the array.

Extending storage using Windows Storage Utilities

Volume extension grows the storage space of a logical drive. During this process, the administrator adds new storage space to an existing logical drive on the same array, usually after the array has been expanded. An administrator may have gained this new storage space by either expansion or by deleting another logical drive on the same array. Unlike drive expansion, the operating system must be aware of changes to the logical drive size.

You extend a volume to:

- Increase raw data storage
- Improve performance by increasing the number of spindles in a logical drive volume
- Change fault-tolerance (RAID) configurations

For more information about RAID levels, see the *Smart Array Controller User Guide*, or the document titled *Assessing RAID ADG vs. RAID 5 vs. RAID 1+0*. Both are available at the Smart Array controller web page or at http://h18000.www1.hp.com/products/servers/proliantstorage/arraycontrollers/documentation.html.

Extend volumes using Disk Management

The Disk Management snap-in provides management of hard disks, volumes or partitions. It can be used to extend a dynamic volume only.

NOTE:

Disk Management cannot be used to extend basic disk partitions.

Guidelines for extending a dynamic volume:

- Use the Disk Management utility.
- You can extend a volume only if it does not have a file system or if it is formatted NTFS.
- You cannot extend volumes formatted using FAT or FAT32.
- You cannot extend striped volumes, mirrored volumes, or RAID 5 volumes.

For more information, see the Disk Management online help.

Expanding storage for EVA arrays using Command View EVA

Presenting a virtual disk offers its storage to a host. To make a virtual disk available to a host, you must present it. You can present a virtual disk to a host during or after virtual disk creation. The virtual disk must be completely created before the host presentation can occur. If you choose host presentation during virtual disk creation, the management agent cannot complete any other task until that virtual disk is created and presented. Therefore, HP recommends that you wait until a virtual disk is created before presenting it to a host.

For more information, see the *HP StorageWorks Command View EVA User Guide*.

Expanding storage using the Array Configuration Utility

The Array Configuration Utility enables online capacity expansion of the array and logical drive for specific MSA storage arrays, such as the MSA1000 and MSA1500. For more information, use the ACU online help, or the procedures to “Expand Array” in the *HP Array Configuration Utility User Guide*

Expand logical drive

This option in the ACU increases the storage capacity of a logical drive by adding unused space on an array to the logical drive on the same array. The unused space is obtained either by expanding an array or by deleting another logical drive on the same array. For more information, use the ACU online help, or the “Extend logical drive” procedure in the *HP Array Configuration Utility User Guide*

Volume shadow copies

NOTE:

Select storage servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses using shadow copies in a non-clustered environment.

The Volume Shadow Copy Service provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. Shadow Copy supports 64 shadow copies per volume.

A shadow copy contains previous versions of the files or folders contained on a volume at a specific point in time. While the shadow copy mechanism is managed at the server, previous versions of files and folders are only available over the network from clients, and are seen on a per folder or file level, and not as an entire volume.

The shadow copy feature uses data blocks. As changes are made to the file system, the Shadow Copy Service copies the original blocks to a special cache file to maintain a consistent view of the file at a particular point in time. Because the snapshot only contains a subset of the original blocks, the cache file is typically smaller than the original volume. In the snapshot's original form, it takes up no space because blocks are not moved until an update to the disk occurs.

By using shadow copies, a storage server can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer. Accessing previous versions of files, or shadow copies, enables users to:

- Recover files that were accidentally deleted. Previous versions can be opened and copied to a safe location.
- Recover from accidentally overwriting a file. A previous version of that file can be accessed.
- Compare several versions of a file while working. Use previous versions to compare changes between two versions of a file.

Shadow copies cannot replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. Because a snapshot only contains a portion of the original data blocks, shadow copies cannot protect against data loss due to media failures. However, the strength of snapshots is the ability to instantly recover data from shadow copies, reducing the number of times needed to restore data from tape.

Shadow copy planning

Before setup is initiated on the server and the client interface is made available to end users, consider the following:

- From what volume will shadow copies be taken?
- How much disk space should be allocated for shadow copies?
- Will separate disks be used to store shadow copies?
- How frequently will shadow copies be made?

Identifying the volume

Shadow copies are taken for a complete volume, but not for a specific directory. Shadow copies work best when the server stores user files, such as documents, spreadsheets, presentations, graphics, or database files.

NOTE:

Shadow copies should not be used to provide access to previous versions of application or e-mail databases.

Shadow copies are designed for volumes that store user data such as home directories and My Documents folders that are redirected by using Group Policy or other shared folders in which users store data.

Shadow copies work with compressed or encrypted files and retain whatever permissions were set on the files when the shadow copies were taken. For example, if a user is denied permission to read a file, that user would not be able to restore a previous version of the file, or be able to read the file after it has been restored.

Although shadow copies are taken for an entire volume, users must use shared folders to access shadow copies. Administrators on the local server must also specify the `\\servername\sharename` path to access shadow copies. If administrators or end users want to access a previous version of a file that does not reside in a shared folder, the administrator must first share the folder.

NOTE:

Shadow copies are available only on NTFS, not FAT or FAT32 volumes.

Files or folders that are recorded by using Shadow Copy appear static, even though the original data is changing.

Allocating disk space

When determining the amount of space to allocate for storing shadow copies, consider both the number and size of files that are being copied, as well as the frequency of changes between copies. For example, 100 files that only change monthly require less storage space than 10 files that change daily. If the frequency of changes to each file is greater than the amount of space allocated to storing shadow copies, no shadow copy is created.

Administrators should also consider user expectations of how many versions they will want to have available. End users might expect only a single shadow copy to be available, or they might expect three days or three weeks worth of shadow copies. The more shadow copies users expect, the more storage space administrators must allocate for storing them.

Setting the limit too low also affects backup programs that use shadow copy technology because these programs are also limited to using the amount of disk space specified by administrators.

NOTE:

Regardless of the volume space that is allocated for shadow copies, there is a maximum of 64 shadow copies for any volume. When the 65th shadow copy is taken, the oldest shadow copy is purged.

The minimum amount of storage space that can be specified is 350 megabytes (MB). The default storage size is 10 percent of the source volume (the volume being copied). If the shadow copies are stored on a separate volume, change the default to reflect the space available on the *storage* volume instead of the *source* volume. Remember that when the storage limit is reached, older versions of the shadow copies are deleted and cannot be restored.

△ CAUTION:

To change the storage volume, shadow copies must be deleted. The existing file change history that is kept on the original storage volume is lost. To avoid this problem, verify that the storage volume that is initially selected is large enough.

Identifying the storage area

To store the shadow copies of another volume on the same file server, a volume can be dedicated on separate disks. For example, if user files are stored on `H:\`, another volume such as `S:\`can be used to store the shadow copies. Using a separate volume on separate disks provides better performance and is recommended for heavily used storage servers.

If a separate volume will be used for the storage area (where shadow copies are stored), the maximum size should be changed to **No Limit** to reflect the space available on the storage area volume instead of the source volume (where the user files are stored).

Disk space for shadow copies can be allocated on either the same volume as the source files or a different volume. There is a trade-off between ease of use and maintenance versus performance and reliability that the system administrator must consider.

By keeping the shadow copy on the same volume, there is a potential gain in ease of setup and maintenance; however, there may be a reduction in performance and reliability.

△ CAUTION:

If shadow copies are stored on the same volume as the user files, note that a burst of disk input/output (I/O) can cause all shadow copies to be deleted. If the sudden deletion of shadow copies is unacceptable to administrators or end users, it is best to use a separate volume on separate disks to store shadow copies.

Determining creation frequency

The more frequently shadow copies are created, the more likely that end users will get the version that they want. However, with a maximum of 64 shadow copies per volume, there is a trade-off between the frequency of making shadow copies and the amount of time that the earlier files will be available.

By default, the storage server creates shadow copies at 0700 and 1200, Monday through Friday. However, these settings are easily modified by the administrator so that the shadow copy schedule can better accommodate end user needs.

Shadow copies and drive defragmentation

When running Disk Defragmenter on a volume with shadow copies activated, all or some of the shadow copies may be lost, starting with the oldest shadow copies.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger. Using this allocation unit size reduces the number of copy outs occurring on the snapshot. Otherwise, the number of changes caused by the defragmentation process can cause shadow copies to be deleted faster than expected. Note, however, that NTFS compression is supported only if the cluster size is 4 KB or smaller.

NOTE:

To check the cluster size of a volume, use the `fsutil fsinfo ntfsinfo` command. To change the cluster size on a volume that contains data, back up the data on the volume, reformat it using the new cluster size, and then restore the data.

Mounted drives

A mounted drive is a local volume attached to an empty folder (called a mount point) on an NTFS volume. When enabling shadow copies on a volume that contains mounted drives, the mounted drives are not included when shadow copies are taken. In addition, if a mounted drive is shared and shadow copies are enabled on it, users cannot access the shadow copies if they traverse from the host volume (where the mount point is stored) to the mounted drive.

For example, assume there is a folder `F:\data\users`, and the `Users` folder is a mount point for `G:\`. If shadow copies are enabled on both `F:\` and `G:\`, `F:\data` is shared as `\\server1\data`, and `G:\data\users` is shared as `\\server1\users`. In this example, users can access previous versions of `\\server1\data` and `\\server1\users` but not `\\server1\data\users`.

Managing shadow copies

The vssadmin tool provides a command line capability to create, list, resize, and delete volume shadow copies.

The system administrator can make shadow copies available to end users through a feature called "Shadow Copies for Shared Folders." The administrator uses the Properties menu (see Figure 5) to turn on the Shadow Copies feature, select the volumes to be copied, and determine the frequency with which shadow copies are made.

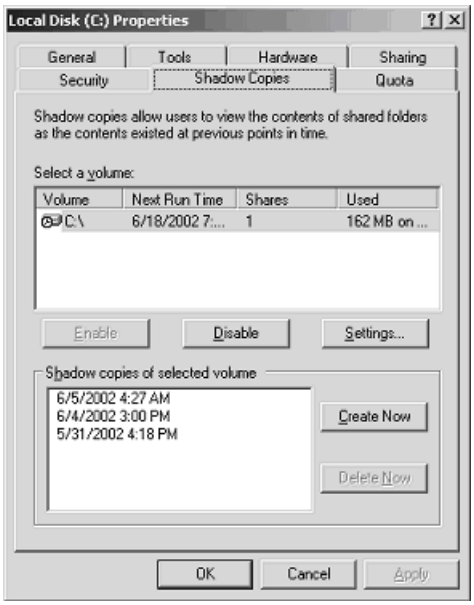

Figure 5 System administrator view of Shadow Copies for Shared Folders

The shadow copy cache file

The default shadow copy settings allocate 10 percent of the source volume being copied (with a minimum of 350 MB), and store the shadow copies on the same volume as the original volume. (See Figure 6). The cache file is located in a hidden protected directory titled "System Volume Information" off of the root of each volume for which shadow copy is enabled.

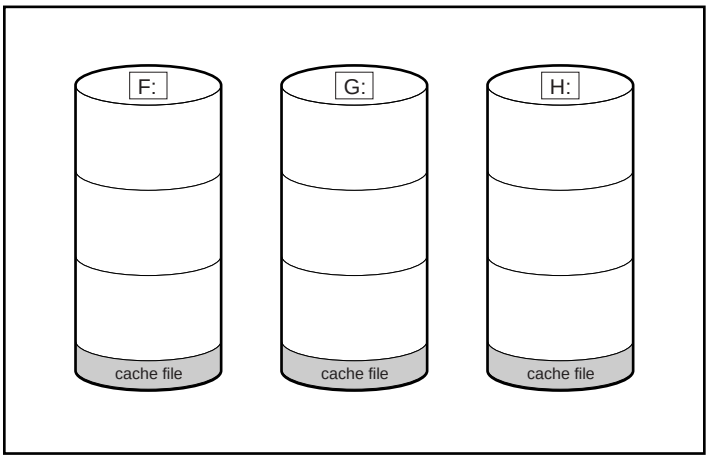

Figure 6 Shadow copies stored on a source volume

The cache file location can be altered to reside on a dedicated volume separate from the volumes containing files shares. (See Figure 7).

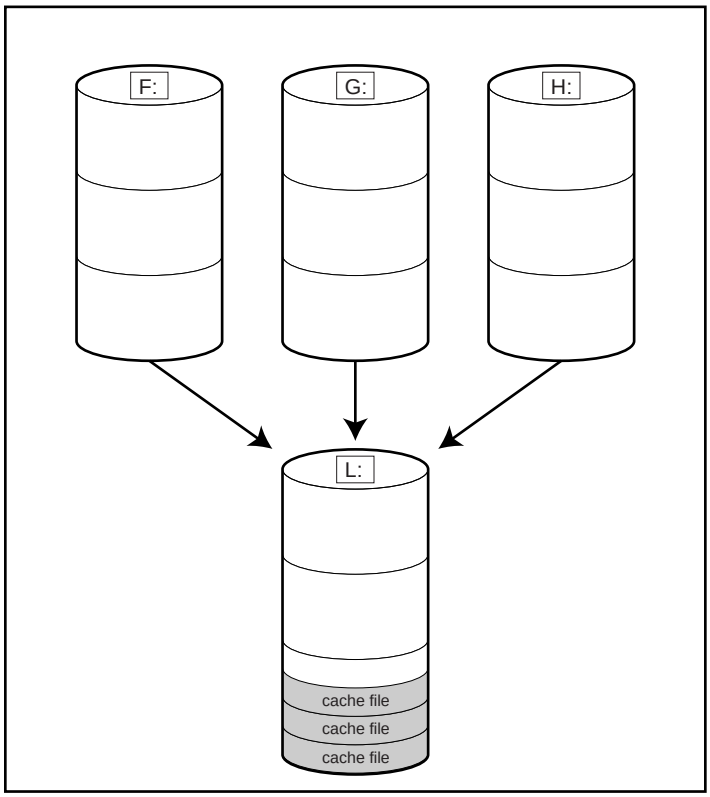

Figure 7 Shadow copies stored on a separate volume

The main advantage to storing shadow copies on a separate volume is ease of management and performance. Shadow copies on a source volume must be continually monitored and can consume space designated for file sharing. Setting the limit too high takes up valuable storage space. Setting the limit too low can cause shadow copies to be purged too soon, or not created at all. By storing shadow copies on a separate volume space, limits can generally be set higher, or set to No Limit. See the online help for instructions on altering the cache file location.

△ CAUTION:

If the data on the separate volume L: is lost, the shadow copies cannot be recovered.

Enabling and creating shadow copies

Enabling shadow copies on a volume automatically results in several actions:

- Creates a shadow copy of the selected volume.
- Sets the maximum storage space for the shadow copies.
- Schedules shadow copies to be made at 7 a.m. and 12 noon on weekdays.

NOTE:

Creating a shadow copy only makes one copy of the volume; it does not create a schedule.

NOTE:

After the first shadow copy is created, it cannot be relocated. Relocate the cache file by altering the cache file location under Properties prior to enabling shadow copy. See "Viewing shadow copy properties" on page 37.

Viewing a list of shadow copies

To view a list of shadow copies on a volume:

1. Access Disk Management.
2. Select the volume or logical drive, then right-click on it.
3. Select **Properties**.
4. Select **Shadow Copies** tab.

All shadow copies are listed, sorted by the date and time they were created.

NOTE:

It is also possible to create new shadow copies or delete shadow copies from this page.

Set schedules

Shadow copy schedules control how frequently shadow copies of a volume are made. There are a number of factors that can help determine the most effective shadow copy schedule for an organization. These include the work habits and locations of the users. For example, if users do not all live in the same time zone, or they work on different schedules, it is possible to adjust the daily shadow copy schedule to allow for these differences.

Do not schedule shadow copies more frequently than once per hour.

NOTE:

When deleting a shadow copy schedule, that action has no effect on existing shadow copies.

Viewing shadow copy properties

The Shadow Copy Properties page lists the number of copies, the date and time the most recent shadow copy was made, and the maximum size setting.

NOTE:

For volumes where shadow copies do not exist currently, it is possible to change the location of the cache file. Managing the cache files on a separate disk is recommended.

CAUTION:

Use caution when reducing the size limit for all shadow copies. When the size is set to less than the total size currently used for all shadow copies, enough shadow copies are deleted to reduce the total size to the new limit. A shadow copy cannot be recovered after it has been deleted.

Redirecting shadow copies to an alternate volume

IMPORTANT:

Shadow copies must be initially disabled on the volume before redirecting to an alternate volume. If shadow copies are enabled and you disable them, a message appears informing you that all existing shadow copies on the volume will be permanently deleted.

To redirect shadow copies to an alternate volume:

1. Access Disk Management.
2. Select the volume or logical drive, then right-click on it.
3. Select **Properties**.
4. Select the **Shadow Copies** tab.
5. Select the volume that you want to redirect shadow copies from and ensure that shadow copies are disabled on that volume; if enabled, click **Disable**.
6. Click **Settings**.
7. In the **Located on this volume** field, select an available alternate volume from the list.

NOTE:

To change the default shadow copy schedule settings, click **Schedule**.

8. Click **OK**.
9. On the **Shadow Copies** tab, ensure that the volume is selected, and then click **Enable**.

Shadow copies are now scheduled to be made on the alternate volume.

Disabling shadow copies

When shadow copies are disabled on a volume, all existing shadow copies on the volume are deleted as well as the schedule for making new shadow copies.

CAUTION:

When the Shadow Copies Service is disabled, all shadow copies on the selected volumes are deleted. Once deleted, shadow copies cannot be restored.

Managing shadow copies from the storage server desktop

To access shadow copies from the storage server desktop:

The storage server desktop can be accessed by using Remote Desktop to manage shadow copies.

1. On the storage server desktop, double-click **My Computer**.
2. Right-click the volume name, and select **Properties**.
3. Click the **Shadow Copies** tab. See Figure 8.

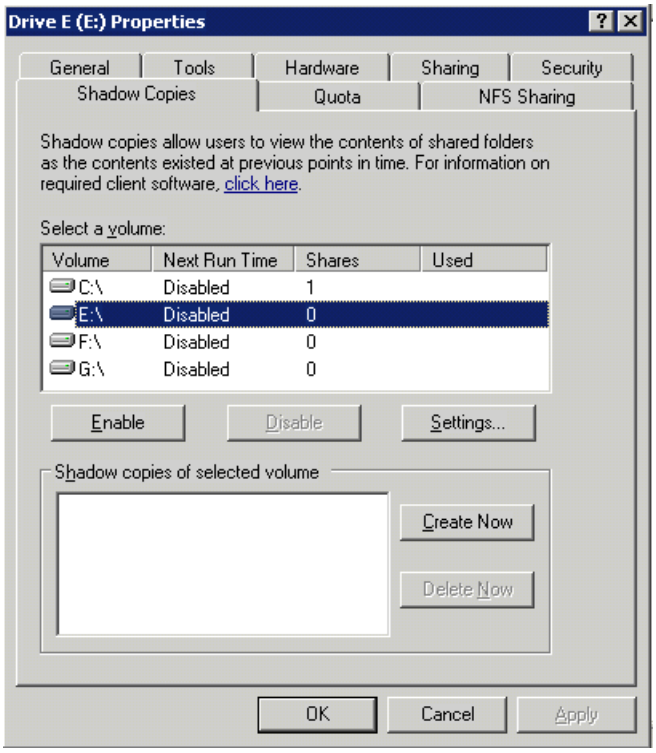

Figure 8 Accessing shadow copies from My Computer

Shadow Copies for Shared Folders

Shadow copies are accessed over the network by supported clients and protocols. There are two sets of supported protocols, SMB and NFS. All other protocols are not supported; this includes HTTP, FTP, AppleTalk, and NetWare Shares. For SMB support, a client-side application denoted as Shadow Copies for Shared Folders is required. The client-side application is currently only available for Windows XP and Windows 2000 SP3+.

No additional software is required to enable UNIX users to independently retrieve previous versions of files stored on NFS shares.

NOTE:

Shadow Copies for Shared Folders supports retrieval only of shadow copies of network shares. It does not support retrieval of shadow copies of local folders.

NOTE:

Shadow Copies for Shared Folders clients are not available for HTTP, FTP, AppleTalk, or NetWare shares. Consequently, users of these protocols cannot use Shadow Copies for Shared Folders to independently retrieve previous versions of their files. However, administrators can take advantage of Shadow Copies for Shared Folders to restore files for these users.

SMB shadow copies

Windows users can independently access previous versions of files stored on SMB shares by using the Shadow Copies for Shared Folders client. After the Shadow Copies for Shared Folders client is installed on the user's computer, the user can access shadow copies for a share by right-clicking on the share to open its Properties window, clicking the **Previous Versions** tab, and then selecting the desired shadow copy. Users can view, copy, and restore all available shadow copies.

Shadow Copies for Shared Folders preserves the permissions set in the access control list (ACL) of the original folders and files. Consequently, users can only access shadow copies for shares to which they have access. In other words, if a user does not have access to a share, he also does not have access to the share's shadow copies.

The Shadow Copies for Shared Folders client pack installs a **Previous Versions** tab in the **Properties** window of files and folders on network shares.

Users access shadow copies with Windows Explorer by selecting **View**, **Copy**, or **Restore** from the **Previous Versions** tab. (See Figure 9). Both individual files and folders can be restored.

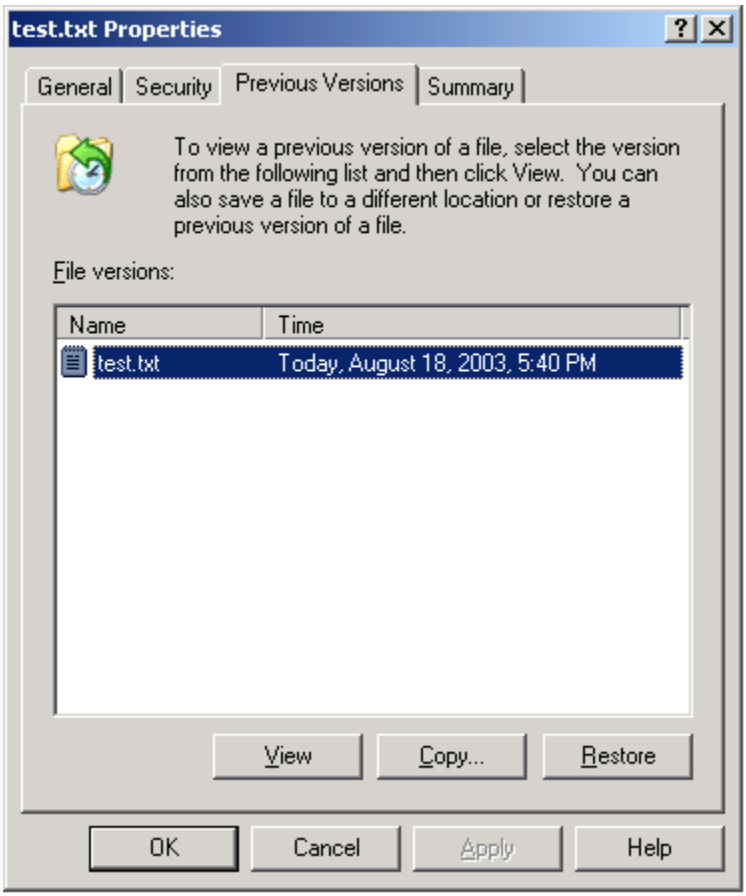

Figure 9 Client GUI

When users view a network folder hosted on the storage server for which shadow copies are enabled, old versions (prior to the snapshot) of a file or directory are available. Viewing the properties of the file or folder presents users with the folder or file history—a list of read-only, point-in-time copies of the file or folder contents that users can then open and explore like any other file or folder. Users can view files in the folder history, copy files from the folder history, and so on.

NFS shadow copies

UNIX users can independently access previous versions of files stored on NFS shares via the NFS client; no additional software is required. Server for NFS exposes each of a share's available shadow copies as a pseudo-subdirectory of the share. Each of these pseudo-subdirectories is displayed in exactly the same way as a regular subdirectory is displayed.

The name of each pseudo-subdirectory reflects the creation time of the shadow copy, using the format .@GMT-YYYY.MM.DD-HH:MM:SS. To prevent common tools from needlessly enumerating the pseudo-subdirectories, the name of each pseudo-subdirectory begins with the dot character, thus rendering it hidden.

The following example shows an NFS share named “NFSShare” with three shadow copies, taken on April 27, 28, and 29 of 2003 at 4 a.m.

NFSShare

.@GMT-2003.04.27-04:00:00

.@GMT-2003.04.28-04:00:00

.@GMT-2003.04.29-04:00:00

Access to NFS shadow copy pseudo-subdirectories is governed by normal access-control mechanisms using the permissions stored in the file system. Users can access only those shadow copies to which they have read access at the time the shadow copy is taken. To prevent users from modifying shadow copies, all pseudo-subdirectories are marked read-only, regardless of the user's ownership or access rights, or the permissions set on the original files.

Server for NFS periodically polls the system for the arrival or removal of shadow copies and updates the root directory view accordingly. Clients then capture the updated view the next time they issue a directory read on the root of the share.

Recovery of files or folders

There are three common situations that may require recovery of files or folders:

- Accidental file deletion, the most common situation
- Accidental file replacement, which may occur if a user selects Save instead of Save As
- File corruption

It is possible to recover from all of these scenarios by accessing shadow copies. There are separate steps for accessing a file compared to accessing a folder.

Recovering a deleted file or folder

To recover a deleted file or folder within a folder:

1. Access to the folder where the deleted file was stored.
2. Position the cursor over a blank space in the folder. If the cursor hovers over a file, that file is selected.
3. Right-click, select **Properties** from the bottom of the menu, and then click the **Previous Versions** tab.
4. Select the version of the folder that contains the file before it was deleted, and then click **View**.
5. View the folder and select the file or folder to recover. The view may be navigated multiple folders deep.
6. Click **Restore** to restore the file or folder to its original location. Click **Copy...** to allow the placement of the file or folder to a new location.

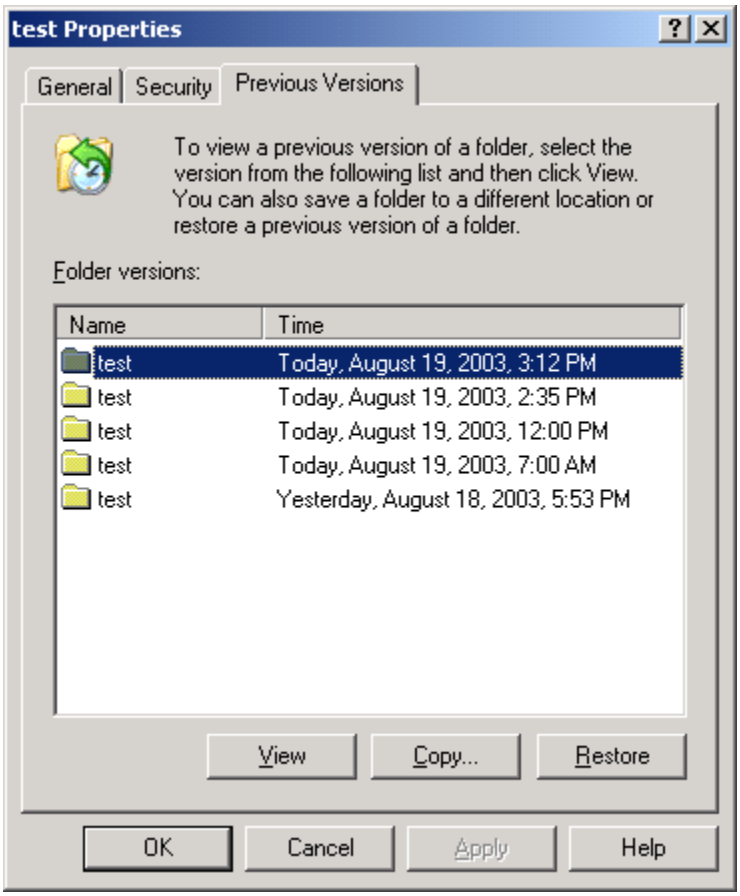

Figure 10 Recovering a deleted file or folder

Recovering an overwritten or corrupted file

Recovering an overwritten or corrupted file is easier than recovering a deleted file because the file itself can be right-clicked instead of the folder. To recover an overwritten or corrupted file:

1. Right-click the overwritten or corrupted file, and then click **Properties**.
2. Click **Previous Versions**.
3. To view the old version, click **View**. To copy the old version to another location, click **Copy...** to replace the current version with the older version, click **Restore**.

Recovering a folder

To recover a folder:

1. Position the cursor so that it is over a blank space in the folder to be recovered. If the cursor hovers over a file, that file is selected.
2. Right-click, select **Properties** from the bottom of the menu, and then click the **Previous Versions** tab.
3. Click either **Copy...** or **Restore**.

   Clicking **Restore** enables the user to recover everything in that folder as well as all subfolders. Clicking **Restore** does not delete any files.

Backup and shadow copies

Shadow copies are only available on the network via the client application, and only at a file or folder level as opposed to the entire volume. Hence, the standard backup associated with a volume backup will not work to back up the previous versions of the file system. To answer this particular issue, shadow copies are available for backup in two situations. If the backup software in question supports the use of shadow copies and can communicate with underlying block device, it is supported, and the previous version of the file system will be listed in the backup application as a complete file system snapshot. If the built-in backup application NTbackup is used, the backup software forces a snapshot, and then uses the snapshot as the means for backup. The user is unaware of this activity and it is not self-evident although it does address the issue of open files.

Shadow Copy Transport

Shadow Copy Transport provides the ability to transport data on a Storage Area Network (SAN). With a storage array and a VSS-aware hardware provider, it is possible to create a shadow copy on one server and import it on another server. This process, essentially "virtual" transport, is accomplished in a matter of minutes, regardless of the size of the data.

NOTE:

Shadow copy transport is supported only on Windows Server 2003 Enterprise Edition, Windows Storage Server 2003 Enterprise Edition, and Windows Server 2003 Datacenter Edition. It is an advanced solution that works only if it has a hardware provider on the storage array.

A shadow copy transport can be used for a number of purposes, including:

- Tape backups

  An alternative to traditional backup to tape processes is transport of shadow copies from the production server onto a backup server, where they can then be backed up to tape. Like the other two alternatives, this option removes backup traffic from the production server. While some backup applications might be designed with the hardware provider software that enables transport, others are not. The administrator should determine whether or not this functionality is included in the backup application.

- Data mining

  The data in use by a particular production server is often useful to different groups or departments within an organization. Rather than add additional traffic to the production server, a shadow copy of the data can be made available through transport to another server. The shadow copy can then be processed for different purposes, without any performance impact on the original server.

The transport process is accomplished through a series of DISKRAID command steps:

1. Create a shadow copy of the source data on the source server (read-only).
2. Mask off (hide) the shadow copy from the source server.
3. Unmask the shadow copy to a target server.
4. Optionally, clear the read-only flags on the shadow copy.

The data is now ready to use.

Folder and share management

The HP ProLiant Storage Server supports several file-sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. This section discusses overview information as well as procedures for the setup and management of the file shares for the supported protocols. Security at the file level and at the share level is also discussed.

NOTE:

Detailed information on setting up and managing NFS and NCP shares is discussed in Microsoft Services for Network File System (MSNFS).

NOTE:

Select servers can be deployed in a clustered or non-clustered configuration. This section discusses share setup for a non-clustered deployment.

Folder management

Volumes and folders on any system are used to organize data. Regardless of system size, systematic structuring and naming conventions of volumes and folders eases the administrative burden. Moving from volumes to folders to shares increases the level of granularity of the types of data stored in the unit and the level of security access allowed.

Folders can be managed using the HP Storage Server Management Console. Tasks include:

- Accessing a specific volume or folder
- Creating a new folder
- Deleting a folder
- Modifying folder properties
- Creating a new share for a volume or folder
- Managing shares for a volume or folder

Managing file-level permissions

Security at the file level is managed using Windows Explorer.

File level security includes settings for permissions, ownership, and auditing for individual files.

To enter file permissions:

1. Using Windows Explorer, access the folder or file that needs to be changed, and then right-click the folder.

2. Click **Properties**, and then click the **Security** tab.

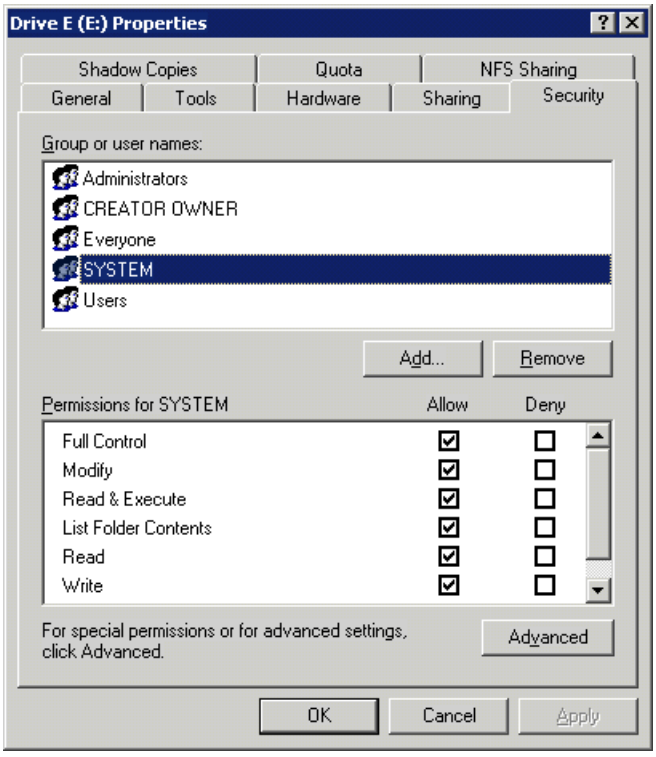

Figure 11 Properties dialog box, Security tab

Several options are available on the **Security** tab:

- To add users and groups to the permissions list, click **Add**. Follow the dialog box instructions.
- To remove users and groups from the permissions list, highlight the desired user or group, and then click **Remove**.
- The center section of the **Security** tab lists permission levels. When new users or groups are added to the permissions list, select the appropriate boxes to configure the common file-access levels.

3. To modify ownership of files, or to modify individual file access level permissions, click **Advanced**.

Figure 12 illustrates the properties available on the **Advanced Security Settings** dialog box.

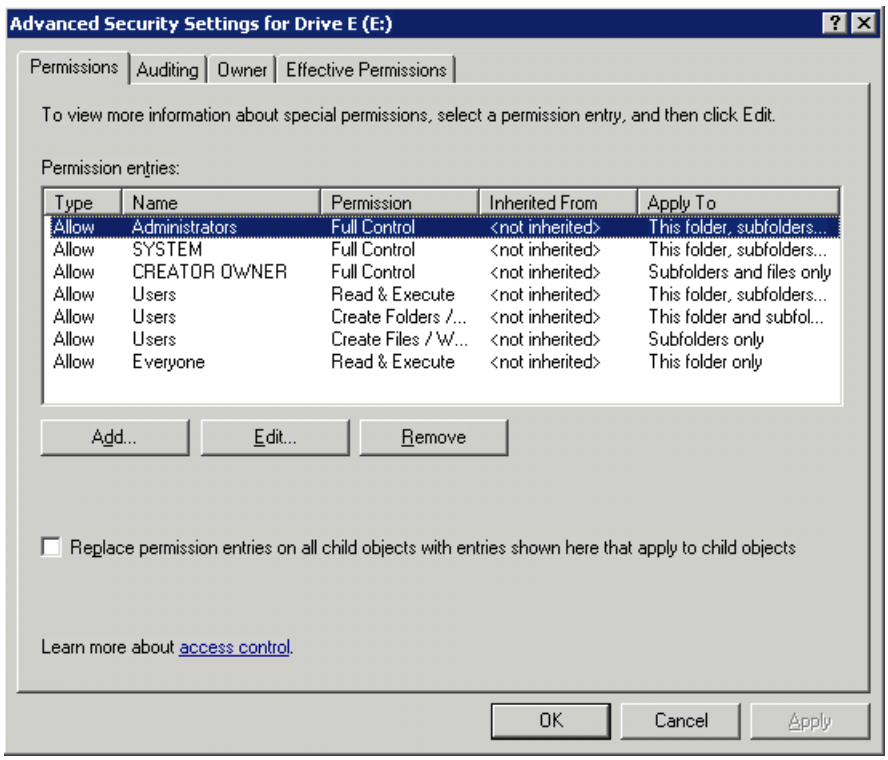

Figure 12 Advanced Security settings dialog box, Permissions tab

Other functionality available in the **Advanced Security Settings** dialog box is illustrated in Figure 12 and includes:

- Add a new user or group—Click **Add**, and then follow the dialog box instructions.
- Remove a user or group— Click **Remove**.
- Replace permission entries on all child objects with entries shown here that apply to child objects—This allows all child folders and files to inherit the current folder permissions by default.
- Modify specific permissions assigned to a particular user or group—Select the desired user or group, and then click **Edit**.

4. Enable or disable permissions by selecting the **Allow** box to enable permission or the **Deny** box to disable permission. If neither box is selected, permission is automatically disabled. Figure 13 illustrates the **Edit** screen and some of the permissions.

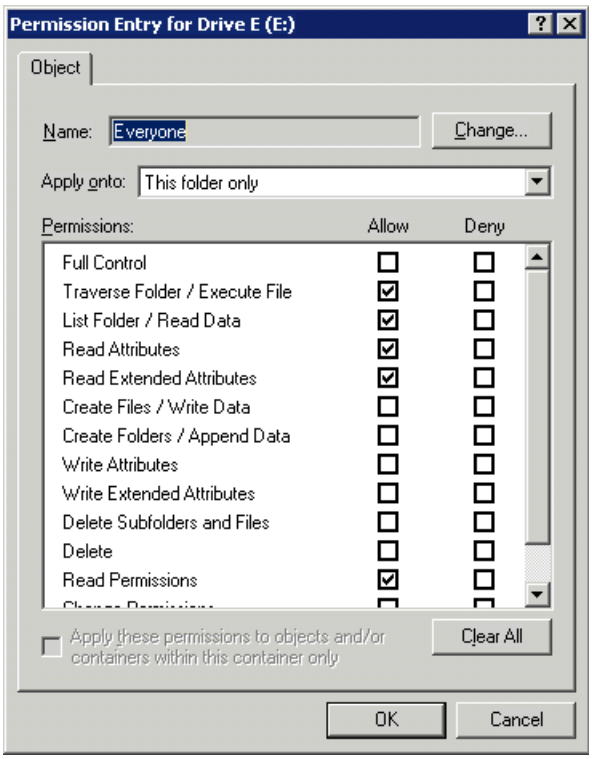

Figure 13 User or group Permission Entry dialog box

Another area of the **Advanced Security Settings** is the **Auditing** tab. Auditing allows you to set rules for the auditing of access, or attempted access, to files or folders. Users or groups can be added, deleted, viewed, or modified through the **Advanced Security Settings Auditing** tab.

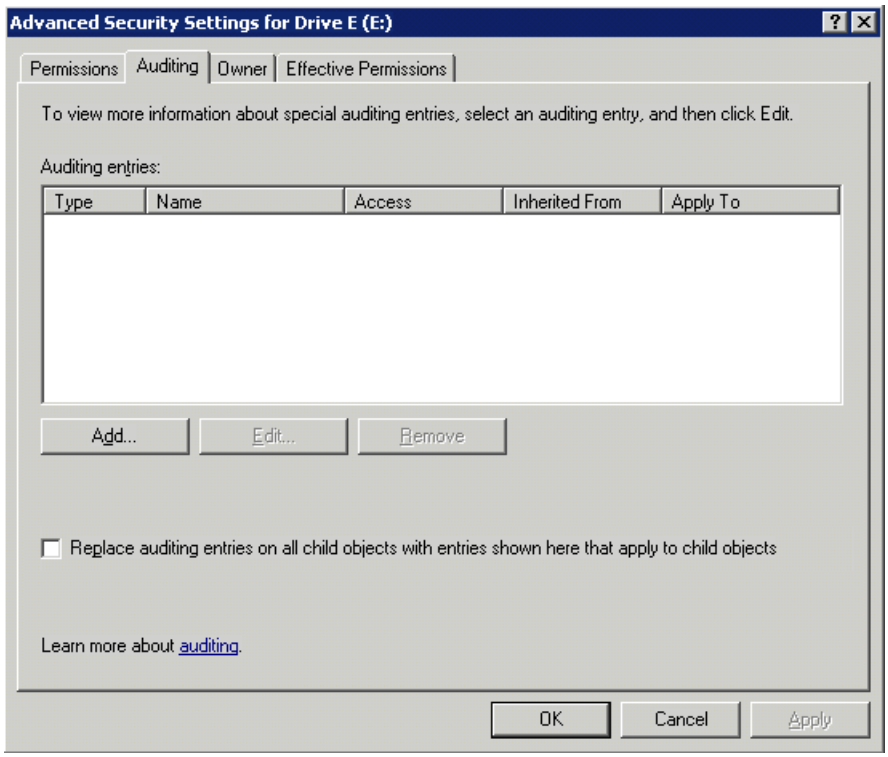

Figure 14 Advanced Security Settings dialog box, Auditing tab

5. Click **Add** to display the Select User or Group dialog box.

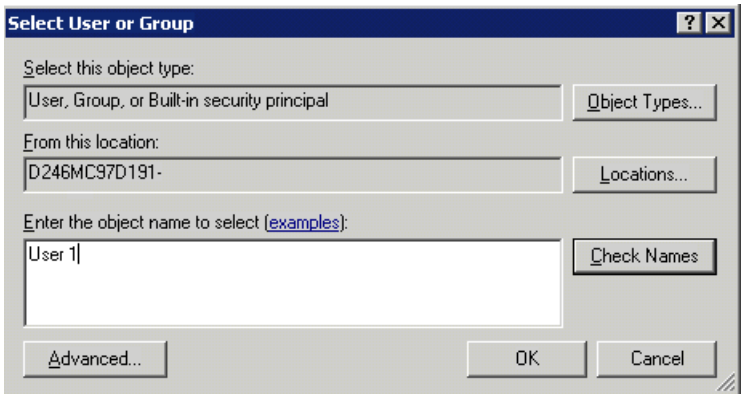

Figure 15 Select User or Group dialog box

NOTE:

Click Advanced to search for users or groups.

6. Select the user or group.

7. Click **OK**.

   The **Auditing Entry** dialog box is displayed.

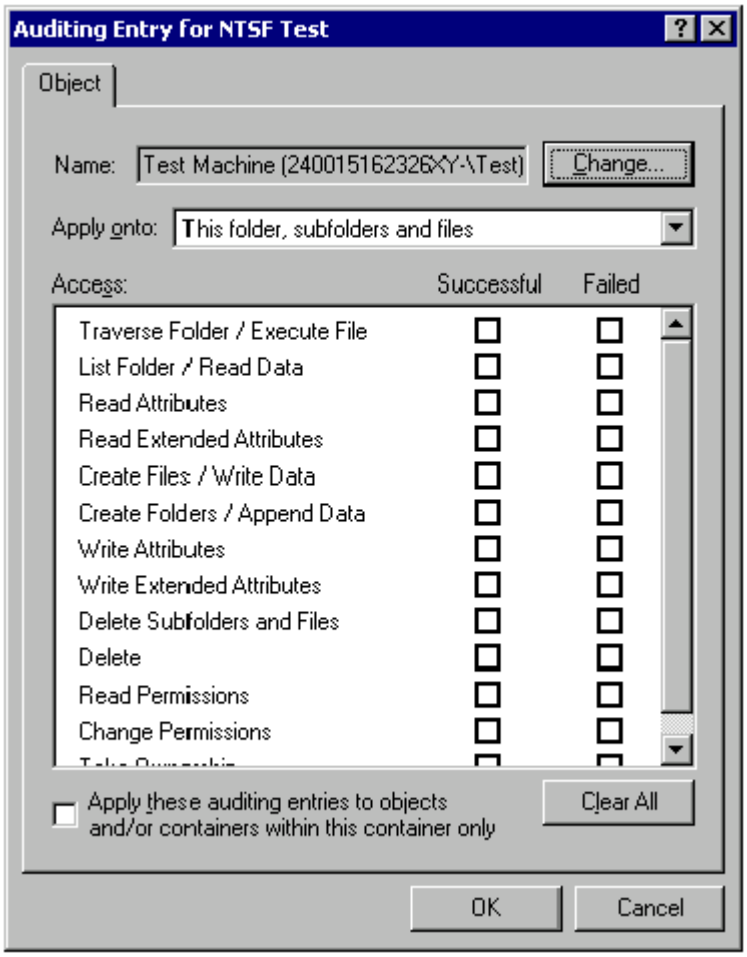

Figure 16 Auditing Entry dialog box for folder name NTFS Test

8. Select the desired **Successful** and **Failed** audits for the user or group.
9. Click **OK**.

NOTE:

Auditing must be enabled to configure this information. Use the local Computer Policy Editor to configure the audit policy on the storage server.

The **Owner** tab allows taking ownership of files. Typically, administrators use this area to take ownership of files when the file ACL is incomplete or corrupt. By taking ownership, you gain access to the files, and then manually apply the appropriate security configurations.

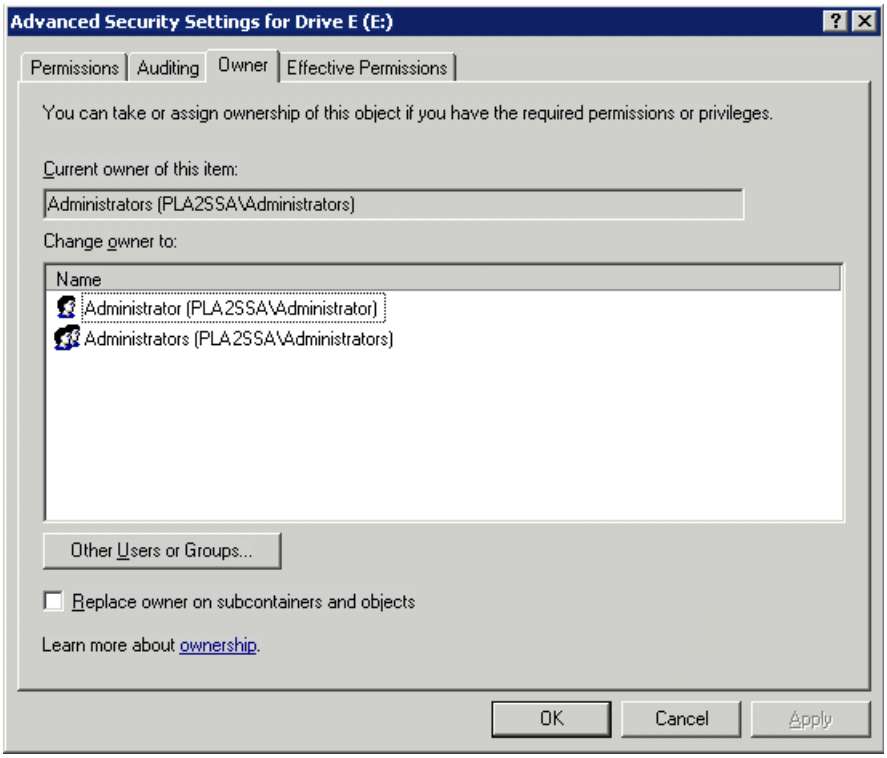

Figure 17 Advanced Security Settings dialog box, Owner tab

The current owner of the file or folder is listed at the top of the screen. To take ownership:

1. Click the appropriate user or group in the **Change owner to** list.
2. If it is also necessary to take ownership of subfolders and files, enable the **Replace owner on subcontainers and objects** box.
3. Click **OK**.

Share management

There are several ways to set up and manage shares. Methods include using Windows Explorer, a command line interface, or the HP Storage Server Management Console.

NOTE:

Select servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses share setup for a non-clustered deployment.

As previously mentioned, the file-sharing security model of the storage server is based on the NTFS file-level security model. Share security seamlessly integrates with file security. In addition to discussing share management, this section discusses share security.

Share considerations

Planning the content, size, and distribution of shares on the storage server can improve performance, manageability, and ease of use.

The content of shares should be carefully chosen to avoid two common pitfalls: either having too many shares of a very specific nature, or of having very few shares of a generic nature. For example, shares for general use are easier to set up in the beginning, but can cause problems later. Frequently, a better approach is to create separate shares with a specific purpose or group of users in mind. However, creating too many shares also has its drawbacks. For example, if it is sufficient to create a single share for user home directories, create a “homes” share rather than creating separate shares for each user.

By keeping the number of shares and other resources low, the performance of the storage server is optimized. For example, instead of sharing out each individual user's home directory as its own share, share out the top-level directory and let the users map personal drives to their own subdirectory.

Defining Access Control Lists

The Access Control List (ACL) contains the information that dictates which users and groups have access to a share, as well as the type of access that is permitted. Each share on an NTFS file system has one ACL with multiple associated user permissions. For example, an ACL can define that User1 has read and write access to a share, User2 has read only access, and User3 has no access to the share. The ACL also includes group access information that applies to every user in a configured group. ACLs are also referred to as permissions.

Integrating local file system security into Windows domain environments

ACLs include properties specific to users and groups from a particular workgroup server or domain environment. In a multidomain environment, user and group permissions from several domains can apply to files stored on the same device. Users and groups local to the storage server can be given access permissions to shares managed by the device. The domain name of the storage server supplies the context in which the user or group is understood. Permission configuration depends on the network and domain infrastructure where the server resides.

File-sharing protocols (except NFS) supply a user and group context for all connections over the network. (NFS supplies a machine-based context.) When new files are created by those users or machines, the appropriate ACLs are applied.

Configuration tools provide the ability to share permissions out to clients. These shared permissions are propagated into a file system ACL, and when new files are created over the network, the user creating the file becomes the file owner. In cases where a specific subdirectory of a share has different permissions from the share itself, the NTFS permissions on the subdirectory apply instead. This method results in a hierarchical security model where the network protocol permissions and the file permissions work together to provide appropriate security for shares on the device.

NOTE:

Share permissions and file-level permissions are implemented separately. It is possible for files on a file system to have different permissions from those applied to a share. When this situation occurs, the file-level permissions override the share permissions.

Comparing administrative (hidden) and standard shares

CIFS supports both administrative shares and standard shares.

- Administrative shares are shares with a last character of $. Administrative shares are not included in the list of shares when a client browses for available shares on a CIFS server.

- Standard shares are shares that do not end in a $ character. Standard shares are listed whenever a CIFS client browses for available shares on a CIFS server.

The storage server supports both administrative and standard CIFS shares. To create an administrative share, end the share name with the $ character when setting up the share. Do not type a $ character at the end of the share name when creating a standard share.

Managing shares

Shares can be managed using the HP Storage Server Management Console. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties
- Publishing in DFS

NOTE:

These functions can operate in a cluster on select servers, but should only be used for non-cluster-aware shares. Use Cluster Administrator to manage shares for a cluster. The page will display cluster share resources.

CAUTION:

Before deleting a share, warn all users to exit that share and confirm that no one is using that share.

File Server Resource Manager

File Server Resource Manager (FSRM) is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. Some of the tasks you can perform are:

- Quota management
- File screening management
- Storage reports

The HP Storage Server Management Console provides access to FSRM tasks.

For procedures and methods beyond what are described below, see the online help. In addition, see a Microsoft File Server Resource Manager white paper available at http://download.microsoft.com/download/7/4/7/7472bf9b-3023-48b7-87be-d2cedc38f15a/WS03R2_Storage_Management.doc .

Quota management

On the Quota Management node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and generate notifications when the quota limits are approached or exceeded.
- Generate auto quotas that apply to all existing folders in a volume or folder, as well as to any new subfolders created in the future.

- Define quota templates that can be easily applied to new volumes or folders and that can be used across an organization.

File screening management

On the File Screening Management node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Create file screens to control the types of files that users can save and to send notifications when users attempt to save blocked files.
- Define file screening templates that can be easily applied to new volumes or folders and that can be used across an organization.
- Create file screening exceptions that extend the flexibility of the file screening rules.

Storage reports

On the Storage Reports node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Schedule periodic storage reports that allow you to identify trends in disk usage.
- Monitor attempts to save unauthorized files for all users or a selected group of users.
- Generate storage reports instantly.

Other Windows disk and data management tools

When you install certain tools, such as Windows Support Tools or Windows Resource Kit Tools, information about these tools might appear in Help and Support Center. To see the tools that are available to you, look in the Help and Support Center under **Support Tasks**, click **Tools**, and then click **Tools by Category**.

NOTE:

The Windows Support Tools and Windows Resource Kit Tools, including documentation for these tools, are available in English only. If you install them on a non-English language operating system or on an operating system with a Multilingual User Interface Pack (MUI), you see English content mixed with non-English content in Help and Support Center. To see the tools that are available to you, click **Start**, click **Help and Support Center**, and then, under **Support Tasks**, click **Tools**.

Additional information and references for file services

Backup

HP recommends that you back up the print server configuration whenever a new printer is added to the network and the print server configuration is modified. For details on implementing the backup solution, see the Medium Business Guide for Backup and Recovery. The guide can be viewed or downloaded from Microsoft at http://www.microsoft.com/technet/itsolutions/smbiz/mits/br/mit_br.mspx.

HP StorageWorks Library and Tape Tools

HP StorageWorks Library and Tape Tools (L&TT) provides functionality for firmware downloads, verification of device operation, maintenance procedures, failure analysis, corrective service actions, and some utility functions. It also provides seamless integration with HP hardware support by generating and e-mailing support tickets that deliver a snapshot of the storage system.

For more information, and to download the utility, see the StorageWorks L&TT web site at http://h18006.www1.hp.com/products/storageworks/ltt.

Antivirus

The server should be secured by installing the appropriate antivirus software. For details on implementing antivirus, see the Medium Business Guide for Antivirus. The guide can be viewed or downloaded from Microsoft at http://www.microsoft.com/technet/itsolutions/smbiz/mits/av/mit_av.mspx.

Security

For guidance on hardening file servers, see the *Microsoft Windows Server 2003 Security Guide*. The guide can be viewed or downloaded at http://www.microsoft.com/technet/security/prodtech/windowsserver2003/w2003hg/sgch00.mspx.

More information

The following web sites provide detailed information for using print services with Windows Server 2003, which also applies to Windows Storage Server 2003.

- Microsoft Storage
  http://www.microsoft.com/windowsserversystem/storage/default.mspx
- Microsoft Windows Storage Server 2003
  http://www.microsoft.com/windowsserversystem/wss2003/default.mspx
- Performance Tuning Guidelines for Windows Server 2003
  http://www.microsoft.com/windowsserver2003/evaluation/performance/tuning.mspx
- Windows SharePoint Services
  http://www.microsoft.com/windowsserver2003/technologies/sharepoint/default.mspx

3 Print services

Microsoft Print Management Console

Print Management in the Microsoft Windows Server 2003 R2 operating system is a Microsoft Management Console (MMC) snap-on that system administrators can use to perform common print management tasks in a large enterprise. It provides a single interface that administrators can use to perform printer and print server management tasks efficiently with detailed control. You can use Print Management from any computer running Windows Server 2003 R2, and you can manage all network printers on print servers running Windows 2000 Server, Windows Server 2003, or Windows Server 2003 R2.

New or improved HP print server features

HP Web Jetadmin

HP Web Jetadmin (WJA)is a web-based tool for remotely installing, configuring, and managing a wide variety of HP and non-HP network peripherals using only a web browser. It supports a modular design whereby plug-ins can be installed to provide additional device, language, and application functionality. WJA is not preinstalled on the storage server, but can be installed (see "HP Web Jetadmin installation" on page 58).

HP Install Network Printer Wizard

The inclusion of the HP Install Network Printer Wizard (INPW) utility on the factory image is new. INPW simplifies the process of installing network printers, including configuration settings on the print server. INPW identifies HP Jetdirect network print devices and allows the user to select the printer to install on the print server.

HP Download Manager for Jetdirect Printer Devices

The inclusion of the HP Download Manager (DLM) for Jetdirect Printer Devices on the factory image is new. DLM is used to upgrade HP Jetdirect print server firmware on HP network printers. The utility obtains the latest firmware catalog from either from the Internet or from a computer with the download firmware images already in place. The DLM discovers all or user-selected Jetdirect devices and upgrades those based on the firmware catalog.

Microsoft Print Migrator utility

The inclusion of the Microsoft Print Migrator utility on the factory image is new. The utility provides complete printer configuration backup of the print server to a user-specified CAB file. Print Migrator supports migration of print configuration data between different versions of Windows, and supports conversion of line printer remote (LPR) ports to the Standard TCP/IP Port Monitor on Windows 2000, Windows XP, and Windows Server 2003.

Network printer drivers

Updated print drivers for HP network printers are preinstalled on the storage server. If a Service Release DVD has been run on the server, there are updated HP network print drivers in the `C:\hpnas\PRINTERS` folder.

Print services management

Print services information to plan, set up, manage, administer, and troubleshoot print servers and print devices are available online using the Help and Support Center feature. To access the Help and Support Center, select **Start > Help and Support**, then **Printers and Faxes** under Help Contents.

Microsoft Print Management Console

The Print Management Console (PMC) can be started from the HP Storage Server Management Console, or the PMC snap-in can be added to the Microsoft Management Console.

HP recommends that you use the *Microsoft Print Management Step-by-Step Guide* on the Documentation CD for print concepts, use of the PMC, and management of network printers. The guide can also be downloaded from http://www.microsoft.com/printserver.

When running the PMC on a server that has Windows Firewall enabled, no printers will be displayed in the printers folder of the PMC. In order for printers to be displayed, you need to open the file and print sharing ports (TCP 139 and 445, and UDP 137 and 138). If this does not fix the problem, or if these ports are already open, you may need to turn off the Windows Firewall to display printers.

To open the file and print sharing ports:

1. Click **Start**, point to Control Panel, and click **Windows Firewall**.
2. On the Exceptions tab, ensure that the File and Printer Sharing check box is selected and click **OK**.

To turn off Windows Firewall:

1. Click **Start**, point to Control Panel, and click **Windows Firewall**.
2. Select **Off** (not recommended) and click **OK**.

HP Web Jetadmin installation

HP Web Jetadmin is used to manage a fleet of HP and non-HP network printers and other peripherals using a web browser. Although not preinstalled, the Web Jetadmin software is located in the `C:\hpnas\Components\WebJetadmin` folder, and can be installed by running the setup program. Follow the installation wizard and supply a password for the local "Admin" username account and a system name.

For more information about Web Jetadmin and Web Jetadmin plug-ins, see http://www.hp.com/go/webjetadmin. For an article on optimizing performance, go to http://h10010.www1.hp.com/wwpc/pscmisc/vac/us/product_pdfs/weboptim.pdf.

Web-based printer management and Internet printing

Internet printing is enabled by default on the print server. Internet printing consists of two main components:

- Web-based printer management with the ability to administer, connect to, and view printers through a web browser.
- Internet printing enabling users to connect to a printer using the printer's URL.

A Microsoft white paper discussing the uses of both components can be obtained at http://www.microsoft.com/windowsserver2003/techinfo/overview/internetprint.mspx.

Planning considerations for print services

Before configuring the print server, the following checklist of items should be followed:

1. **Determine the operating system version of the clients that will send jobs to this printer**. This information is used to select the correct client printer drivers for the client and server computers using the printer. Enabling this role on the print server allows the automatic distribution of these drivers to the clients. Additionally, the set of client operating systems determines which of these drivers need to be installed on the server during the print server role installation.
2. **At the printer, print a configuration or test page that includes manufacturer, model, language, and installed options**. This information is needed to choose the correct printer driver. The manufacturer and model are usually enough to uniquely identify the printer and its language. However, some printers support multiple languages, and the configuration printout usually lists them. Also, the configuration printout often lists installed options, such as extra memory, paper trays, envelope feeders, and duplex units.
3. **Choose a printer name**. Users running Windows-based client computers choose a printer by using the printer name. The wizard that you will use to configure your print server provides a default name, consisting of the printer manufacturer and model. The printer name is usually fewer than 31 characters in length.
4. **Choose a share name**. A user can connect to a shared printer by entering this name, or by selecting it from a list of share names. The share name is usually fewer than 8 characters for compatibility with MS-DOS and Windows 3.x clients.
5. (Optional) **Choose a location description and a comment**. These can help identify the location of the printer and provide additional information. For example, the location could be "Second floor, copy room" and the comment could be "Additional toner cartridges are available in the supply room on floor 1."
6. **Enable management features for Active Directory and Workgroup Environments**. If the print server is part of an Active Directory domain rather than Workgroup, the print server enables the following management features:
   - Restrict access to printer-based domain user accounts.
   - Publish shared printers to Active Directory to aid in search for the resource.
7. **Deploy printers using group policy**. Print management can be used with Group Policy to automatically add printer connections to a server's Printers and Faxes folder. For more information, see the Microsoft article at http://technet2.microsoft.com/WindowsServer/en/Library/ab8d75f8-9b35-4e3e-a344-90d7799927231033.mspx.
8. **Determine whether printer spooling be enabled**. Two or more identical printers that are connected to one print server can act as a single printer. As a means to load-balance print queues when you print a document, the print job is sent to the first available printer in the pool. See "Setting printer properties" in the Windows online help for additional information.

Print queue creation

In addition to Windows Printer and Faxes, Add Printer Wizard, the HP Install Network Printer Wizard (INPW) utility discovers HP Jetdirect network printers on the local network and allows print queues to

be created on the print server. The utility is located on the storage server in the `C:\hpnas\Components\Install Network Printer Wizard` folder.

Sustaining print administration tasks

Tasks that need to be performed regularly to support the print services include:

- Monitoring print server performance using the built-in performance monitoring tool in the Windows Server operating system.
- Supporting printers that include adding, moving, and removing printers as requirements change.
- Installing new printer drivers.
- Recording information about the printer's name, share names, printer features, and the location where the printers are physically installed. This information should be kept in an easily accessible place.

For process suggestions for recurring tasks, see the Microsoft Print Service Product Operations Guide at http://www.microsoft.com/technet/itsolutions/cits/mo/winsrvmg/pspog/pspog3/mspx.

Driver updates

Print drivers

The latest print drivers for many HP network printers are supplied on the Service Release DVD. If selected as part of the service release installation process, updated print drivers are copied to the print drivers folder `C:\hpnas\PRINTERS` on the storage server. Print drivers are also available for download on the HP Support web site for individual network printers.

User-mode vs. kernel-mode drivers

Drivers can be written in either user mode (also called version 3 drivers) or kernel mode (also called version 2 drivers). Native drivers on Windows 2000 and later run in user mode. Windows Server 2003 and Windows Storage Server 2003 can run kernel-mode drivers, although this is not recommended for stability reasons.

Kernel-mode driver installation blocked by default

In Windows Server 2003 and Windows Storage Server 2003, installation of kernel-mode drivers is blocked by default.

To allow kernal-mode drivers to be installed, perform the following steps:

1. Open Group Policy, click **Start > Run**, then type **gpedit.msc**, and press **Enter**.
2. Under **Local Computer Policy**, double-click **Computer Configuration**.
3. Right-click **Disallow installation of printers using kernel-mode drivers**, and then click **Properties**.
4. On the **Setting** tab, click either **Not Configured** or **Disabled**, and then click **OK**.

HP Jetdirect firmware

The HP Download Manager (DLM) utility for Jetdirect printers provides upgrades of HP Jetdirect print server firmware on HP network printers. The utility is located on the storage server in the `C:\hpnas\Components\Download Manager for Jetdirect` folder. A connection to the

Internet is required, or the utility can be pointed to a local location where the firmware images are stored. For more information on upgrading HP Jetdirect print server firmware, see http://h20000.www2.hp.com/bizsupport/TechSupport/Document.jsp?objectID=bpj06917.

Printer server scalability and sizing

A Microsoft technical paper overviews several key factors that influence the capacity of a given print server configuration. While this paper cannot provide a predictive formula to determine the printing throughput of a given configuration, it does describe several reference systems and their capacity. This paper also presents the information necessary to help the system administrator or capacity planner estimate, and later monitor, their server workload. The current version of this paper is maintained at http://www.microsoft.com/printserver.

Backup

It is recommended that you back up the print server configuration whenever a new printer is added to the network and the print server configuration is modified. For details on implementing the backup solution, see the *Medium Business Guide for Backup and Recovery*. The guide can be viewed or downloaded from Microsoft at http://www.microsoft.com/technet/itsolutions/smbiz/mits/br/mit_br.mspx.

The Print Migrator utility is recommended as a print-specific alternative to backing up print configuration settings on the print server. The Print Migrator utility is located in the `C:\hpnas\Components\PrintMigrator` folder on the storage server.

For more information about the Print Migrator utility, see http://www.microsoft.com/WindowsServer2003/techinfo/overview/printmigrator3.1.mspx.

Best practices

The following is practical advice for managing print devices:

- Printers and print servers should be published in Active Directory.
- Locate printers in common areas, such as near conference rooms.
- Protect print servers using antivirus software.
- Ensure the print server is included in the backup configuration.
- Use Microsoft Printer Migrator to back up a print server configuration and restore settings on a new print server. This eliminates the need to manually re-create print queues and printer ports, install drivers, and change the IP configuration.
- Use Microsoft Printer Migrator to backup new printers configured on the print server.
- Use Microsoft Printer Migrator when migrating to new print servers.
- Perform a full backup of the print server, including the state information, before releasing the system to the users in the production environment.
- Whenever a new configuration is made or existing configuration is modified, a backup should be performed.
- To optimize performance, move the print spooler to another disk, separate from the disk supporting the operating system. To move the print spooler to another disk:
  - Start Printer and Faxes.
  - On the File menu, click **Server Properties**, and then click the **Advanced** tab.
  - In the Spool folder window, enter the path and the name of the new default spool folder for the print server and then click **Apply** or **OK**.

- Stop and restart the spooler service, or restart the print server.

Troubleshooting

The online help or Help and Support Center feature should be used to troubleshoot general and common print-related problems. Printing help can be accessed by selecting **Start > Help and Support**, then the **Printers and Faxes** selection under **Help Contents**.

The same print troubleshooting information can be accessed at http://www.microsoft.com/technet/prodtechnol/windowsserver2003/library/ServerHelp/2048a7ba-ec57-429c-95a3-226eea32d126.mspx

Specific print server related problems as well as other system related known issues and workarounds are addressed in release notes. To view the latest version, visit http://www.hp.com/go/support. Select **See support and troubleshooting information** and enter a product name/number. Under **self-help resources**, select the **manuals (guides, supplements, addendums, etc)** link.

Additional references for print services

The following web sites provide detailed information for using print services with Windows Server 2003, which also applies to Windows Storage Server 2003.

- Windows Server 2003 print services home page at http://www.microsoft.com/windowsserver2003/technologies/print/default.mspx
- Medium Business Solution for Print Services at http://www.microsoft.com/technet/itsolutions/smbiz/mits/ps/mit_ps.mspx.

4 Microsoft Services for Network File System (MSNFS)

This chapter discusses networking features in Microsoft Services for Network File System (MSNFS).

MSNFS Features

MSNFS is an update to the NFS components that were previously available in Services for UNIX 3.5.

MSNFS includes the following new features:

- Updated administration snap-in—MSNFS Administration
- Active Directory Lookup—The Identity Management for UNIX Active Directory schema extension, available in Microsoft Windows Server 2003 R2, includes UNIX user identifier (UID) and group identifier (GID) fields, which enables Server for NFS and Client for NFS to look up Windows-to-UNIX user account mappings directly from Active Directory. Identity Management for UNIX simplifies Windows-to-UNIX user account mapping management in Active Directory.
- Enhanced server performance—Microsoft Services for NFS includes a file filter driver, which significantly reduces common server file access latencies.
- UNIX special device support—Microsoft Services for NFS supports UNIX special devices (mknod).
- Enhanced UNIX support—Microsoft Services for NFS now supports the following versions of UNIX:
  - Hewlett Packard HP-UX version 11i
  - IBM AIX version 5L 5.2
  - Red Hat Linux version 9
  - Sun Microsystems Solaris version 9

The following features that were previously available in Services for UNIX 3.5 are not included in MSNFS:

- Gateway for NFS
- Server for PCNFS
- All PCNFS components of Client for NFS

UNIX Identity Management

Identity Management for UNIX makes it easy to integrate users of Windows operating systems into existing UNIX environments. It provides manageability components that simplify network administration and account management across both platforms.

With Identity Management for UNIX, the administrator can:

- Manage user accounts and passwords on Windows and UNIX systems using Network Information Service (NIS).
- Automatically synchronize passwords between Windows and UNIX operating systems.

UNIX Identity Management consists of the following components:

- Administration components
- Password synchronization
- Server for NIS

The UNIX Identity Management component is not enabled by default on the storage server. To install this component:

1. Access **Add/Remove Programs**.
2. Select **Add/Remove Windows Components > Active Directory Services > Details**.
3. Install **Identity Management for UNIX**.

MSNFS use scenarios

The following use scenarios are supported by MSNFS file services:

- Allow UNIX clients to access resources on computers running Windows Server 2003 R2.

  Your company may have UNIX clients accessing resources, such as files, on UNIX file servers. To take advantage of new Windows Server 2003 features, such as Shadow Copies for Shared Folders, you can move resources from your UNIX servers to computers running Windows Server 2003 R2. You can then set up MSNFS to enable access by UNIX clients that are running NFS software. All of your UNIX clients will be able to access the resources using the NFS protocol with no changes required.

- Allow computers running Windows Server 2003 R2 to access resources on UNIX file servers.

  Your company may have a mixed Windows and UNIX environment with resources, such as files, stored on UNIX file servers. You can use MSNFS to enable computers running Windows Server 2003 R2 to access these resources when the file servers are running NFS software.

NOTE:

Services for NFS can be implemented in both clustered and non-clustered environments using select storage servers. This chapter discusses Services for NFS in a non-clustered deployment. If your storage server is capable of using clusters, see the Cluster administration chapter for more information. (This chapter is not in manuals for those models that cannot use clusters.)

MSNFS components

MSNFS comprises the following three main components:

- Username Mapping Server

  Username Mapping Server maps user names between Windows and UNIX user accounts. In a heterogeneous network, users have separate Windows and UNIX security accounts. Users must provide a different set of credentials to access files and other resources, depending on whether they are stored on a Windows or UNIX file server. To address this issue, Username Mapping Server maps the Windows and UNIX user names so that users can log on with either their Windows or UNIX credentials and access resources regardless of whether they are stored on a Windows or UNIX file server.

- Server for NFS

  Normally, a UNIX computer cannot access files on a Windows-based computer. A computer running Windows Server 2003 R2 and Server for NFS, however, can act as a file server for both Windows and UNIX computers.

- Client for NFS

  Normally, a Windows-based computer cannot access files on a UNIX computer. A computer running Windows Server 2003 R2 and Client for NFS, however, can access files stored on a UNIX-based NFS server.

The Client for NFS feature of the Microsoft Services for NFS component is not preinstalled on the storage server although information about this feature appears in the online help. To enable Client for NFS:

1. Go to **Add/Remove Programs**.
2. Select **Add/Remove Windows Components > Other Network File and Print Services > Microsoft Services for NFS > Details**.
3. Install Client for NFS.

Administering MSNFS

To access Microsoft Services for Network File System from the Start menu:

1. Select **Start > Programs > Administrative Tools**.
2. Click **Microsoft Services for Network File System**.

To access Microsoft Services for Network File System from the HP Storage Server Management console:

1. Access the HP Storage Server Management console by clicking on the shortcut icon on the desktop.
2. In the left pane of the console, select the **Share Folder Management** listing.
3. In the center pane, under **Share Utilities**, select **Microsoft Services for NFS** (see Figure 18).

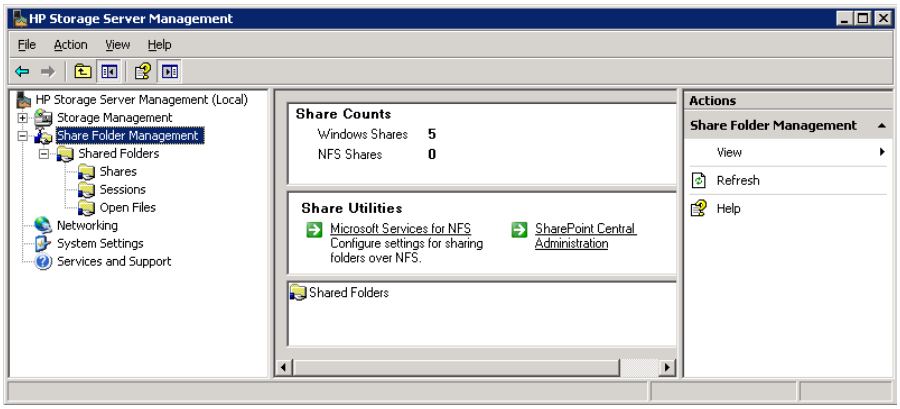

Figure 18 Accessing MSNFS from HP Storage Server Management console

Server for NFS

With Server for NFS, a computer running the Microsoft Windows Server 2003 R2 operating system can act as a Network File System (NFS) server. Users can then share files in a mixed environment of computers, operating systems, and networks. Users on computers running NFS client software can gain access to directories (called shares) on the NFS server by connecting (mounting) those directories to their computers. From the viewpoint of the user on a client computer, the mounted files are indistinguishable from local files.

UNIX computers follow advisory locking for all lock requests. This means that the operating system does not enforce lock semantics on a file, and applications that check for the existence of locks can use these locks effectively. However, Server for NFS implements mandatory locks even for those locking requests that are received through NFS. This ensures that locks acquired through NFS are visible through the server message block (SMB) protocol and to applications accessing the files locally. Mandatory locks are enforced by the operating system.

Server for NFS Authentication DLL versus Service for User for Active Directory domain controllers

On a Windows Storage Server 2003 R2 storage server, Server for NFS depends on a domain controller feature called Service for User (S4U) to authenticate UNIX users as their corresponding Windows users. Windows Server operating systems prior to Windows Server 2003 and Windows Storage Server 2003 do not support S4U. Also, in mixed domain environments, legacy Services for UNIX (SFU), Services for NFS and Windows Storage Server 2003 NFS deployments do not use the S4U feature and still depend on the Server for NFS Authentication DLL being installed on domain controllers.

Therefore, the administrator needs to install the Server for NFS Authentication DLL on Windows 2000 domain controllers when:

- The NFS file serving environment uses previous NFS releases (NAS, SFU, and so on).
- The Windows domain environment uses pre-2003 domain controllers.

See Table 4 for guidance as to when to use NFS Authentication DLL instead of S4U legacy NFS and R2 MSNFS.

Table 4 Authentication table

Domain controller type	Legacy NFS (pre-WSS2003 R2)	MSNFS (WSS2003 R2)
Legacy domain controller (pre-WSS2003)	Requires NFS Authentication DLL on domain controller	Requires NFS Authentication DLL on domain controller
Recent domain controllers (WSS2003 and later)	Requires NFS Authentication DLL on domain controller	Uses the built-in S4U (on the domain controller). It is unaffected by the NFS Authentication DLL on the domain controller.

The S4U set of extensions to the Kerberos protocol consists of the Service-for-User-to-Proxy (S4U2Proxy) extension and the Service-for-User-to-Self (S4U2Self) extension. For more information about the S4U2 extensions, see the Kerberos articles at the following URLs: http://searchwindowssecurity.techtarget.com/originalContent/0,289142,sid45_gci1013484,00.html (intended for IT professionals) and http://msdn.microsoft.com/msdnmag/issues/03/04/SecurityBriefs/default.aspx (intended for developers).

Installing NFS Authentication DLL on domain controllers

NOTE:

If the authentication software is not installed on all domain controllers that have user name mappings, including primary domain controllers, backup domain controllers, and Active Directory domains, then domain user name mappings will not work correctly.

You need to install the version of NFS Authentication included with Services for UNIX 3.5. You can download Services for UNIX 3.5 at no charge from http://go.microsoft.com/fwlink/?LinkId=44501.

To install the Authentication software on the domain controllers:

1. From the SFU 3.5 files, locate the directory named `SFU35SEL_EN`.
2. On the domain controller where the Authentication software is being installed use Windows Explorer to:
   a. Open the shared directory containing `setup.exe`.
   b. Double-click the file to open it. Windows Installer is opened.

NOTE:

If the domain controller used does not have Windows Installer installed, locate the file `InstMSI.exe` on the SFU 3.5 directory and run it. After this installation, the Windows Installer program starts when opening `setup.exe`.

3. In the Microsoft Windows Services for UNIX Setup Wizard dialog box, click **Next**.
4. In the User name box, enter your name. If the name of your organization does not appear in the Organization box, enter the name of your organization there.
5. Read the End User License Agreement carefully. If you accept the terms of the agreement, click **I accept the terms in the License Agreement**, and then click **Next** to continue installation. If you click **I do not accept the License Agreement** (Exit Setup), the installation procedure terminates.
6. Click **Custom Installation**, and then click **Next**.
7. In the Components pane, click the down arrow next to Windows Services for UNIX, and then click **Entire component will not be available**.
8. Click the plus sign (+) next to Authentication Tools.
9. In the Components pane, click the plus sign (+) next to Authentication Tools.
10. Click **Server for NFS Authentication**, click **Will be installed on local hard drive**, and then click **Next**.
11. Follow the remaining instructions in the Wizard.

NOTE:

NFS users can be authenticated using either Windows domain accounts or local accounts on the Windows server. Server for NFS Authentication must be installed on all domain controllers in the domain if NFS users will be authenticated using domain accounts. Server for NFS Authentication is always installed on the computer running Server for NFS.

Elevate S4U2 functionality on Windows Server 2003 domain controllers

NOTE:

The S4U2 functionality does not work until the domain functional level is elevated to Windows Server 2003.

To elevate the functional level to Windows Server 2003:

1. On the Windows 2003 domain controller, open Active Directory Domains and Trusts.

2. In the console tree, right-click the domain for which you want to raise functionality, and then click **Raise Domain Functional Level**.
3. In Select an available domain functional level, click **Windows Server 2003**.
4. Click **Raise**.

Server for NFS administration

The Server for NFS administration online help contains information for the following topics:

- Understanding the Server for NFS component
- Starting and stopping Server for NFS
- Configuring Server for NFS
- Securing Server for NFS
- Optimizing Server for NFS performance
- Using file systems with NFS
- Managing NFS shares
- Managing NFS client groups
- Using Microsoft Services for NFS with server clusters
- Server for NFS Authentication

Accessing NFS resources for Windows users and groups

Server for NFS allows Windows clients to access NFS resources on the storage server without separately logging on to Server for NFS. The first time users attempt to access an NFS resource, the Server for NFS looks up the user's UNIX UID and GID information in either Windows Active Directory or the User Name Mapping function on the storage server. If the UNIX UID and GID information is mapped to a Windows user and group accounts, the Windows names are returned to Server for NFS, which then uses the Windows user and group names to grant file access. If the UNIX UID and GID information is not mapped, then Server for NFS will deny file access.

There are two ways to specify how Server for NFS on the storage server obtains Windows user and group information:

- Using the Windows interface
- Using a command line (`nfsadmin.exe`)

IMPORTANT:

- Before using Active Directory Lookup, administrators must install and populate the Identity Management for UNIX Active Directory schema extension, included in Windows Server 2003 R2, or have an equivalent schema which includes UNIX UID and GID fields.
- The IP address of the User Name Mapping server can be specified instead of the name of the server.
- Before using User Name Mapping, the computer running Server for NFS must be listed in the .maphosts file on the computer running User Name Mapping. For more information, see "Securing access to the User Name Mapping server."

For additional information about accessing NFS resources, see the MSNFS online help. For additional information about Identity Management for UNIX, see the UNIX Identify Management online help

Managing access using the .maphosts file

The User Name Mapping component of MSNFS acts as an intermediary between NFS servers and NFS clients on a network containing UNIX hosts and Windows-based computers. To maintain the implicit trust relationship between NFS client and host computers, administrators can control which computers can access User Name Mapping by editing the .maphosts in the %windir%\msnfs directory of the storage server. Conditions to allow or deny access include:

- If the .maphosts file is present but not empty, then only those computers allowed access by entries in the file can access User Name mapping.
- If the .maphosts file is present but empty (the default), no computers except the computer running User Name Mapping itself can access User Name Mapping.
- If the .maphosts file is not present, no computers (including the computer running User Name Mapping) can access User Name Mapping.

The ordering of entries is important as User Name Mapping searches the .maphosts file from the top down until it finds a match.

For additional information about the .maphosts file, see the MSNFS online help.

Allowing anonymous access to resources by NFS clients

You may want to add anonymous access to a share, for example when it is not desirable or possible to create and map a UNIX account for every Windows user. A UNIX user whose account is not mapped to a Windows account is treated by Server for NFS as an anonymous user. By default, the user identifier (UID) and group identifier (GID) is -2.

For example, if files are created on an NFS Share by UNIX users who are not mapped to Windows users, the owner of those files are listed as anonymous user and anonymous group, (-2,-2).

By default, Server for NFS does not allow anonymous users to access a shared directory. When an NFS share is created, the anonymous access option can be added to the NFS share. The values can be changed from the default anonymous UID and GID values to the UID and GID of any valid UNIX user and group accounts.

NOTE:

In Windows Server 2003, the Everyone group does not include anonymous users by default.

When allowing anonymous access to an NFS Share, the following must be performed by a user with administrative privileges due to Windows Storage Server 2003 security with anonymous users and the Everyone group.

1. Click **Remote Desktop**. Log on to the storage server.
2. Click **Start** >**Control Panel** > **Administrative Tools**, and then click **Local Security Policy**.
3. In Security Settings, double-click **Local Policies**, and then click **Security Options**.
4. Right-click **Network access: Let Everyone permissions apply to anonymous users**, and then click **Properties**.
5. To allow permissions applied to the Everyone group to apply to anonymous users, click **Enabled**. The default is **Disabled**.
6. Restart the NFS server service. From a command prompt, enter `net stop nfssvc`. Then enter `net start nfssvc`. Notify users before restarting the NFS service.

7. Assign the Everyone group the appropriate permissions on the NFS Share.
8. Enable anonymous access to the share.

To enable anonymous access to an NFS share, do the following:

1. Open Windows Explorer by clicking **Start** > **Run**, and entering Explorer.
2. Navigate to the NFS share.
3. Right-click the NFS Share, and then click **Properties**.
4. Click **NFS Sharing**.
5. Select the **Allow Anonymous Access** checkbox.
6. Change from the default of -2,-2, if desired.
7. Click **Apply**.
8. Click **OK**.

Best practices for running Server for NFS

- Provide user-level security.
- Secure files.
- Secure new drives.
- Allow users to disconnect before stopping the Server for NFS service.
- Use naming conventions to identify shares with EUC encoding.
- Protect configuration files.

For further details, see the online help for Microsoft Services for Network File System.

User Name Mapping

The User Name Mapping component provides centralized user mapping services for Server for NFS and Client for NFS. User Name Mapping lets you create maps between Windows and UNIX user and group accounts even though the user and group names in both environments may not be identical. User Name Mapping lets you maintain a single mapping database making it easier to configure account mapping for multiple computers running MSNFS.

In addition to one-to-one mapping between Windows and UNIX user and group accounts, User Name Mapping permits one-to-many mapping. This lets you associate multiple Windows accounts with a single UNIX account. This can be useful, for example, when you do not need to maintain separate UNIX accounts for individuals and would rather use a few accounts to provide different classes of access permission.

You can use simple maps, which map Windows and UNIX accounts with identical names. You can also create advanced maps to associate Windows and UNIX accounts with different names, which you can use in conjunction with simple maps.

User Name Mapping can obtain UNIX user, password, and group information from one or more Network Information Service (NIS) servers or from password and group files located on a local hard drive. The password and group files can be copied from a UNIX host or from a NIS server.

User Name Mapping periodically refreshes its mapping database from the source databases, ensuring that it is always kept up-to-date as changes occur in the Windows and UNIX name spaces. You can also refresh the database anytime you know the source databases have changed.

You can back up and restore User Name Mapping data at any time. Because the database is backed up to a file, you can use that file to copy the mapping database to another server. This provides redundancy for the sake of fault tolerance.

NOTE:

If you obtain information from multiple NIS domains, it is assumed that each domain has unique users and user identifiers (UIDs). User Name Mapping does not perform any checks.

User Name Mapping associates Windows and UNIX user names for Client for NFS and Server for NFS. This allows users to connect to Network File System (NFS) resources without having to log on to UNIX and Windows systems separately.

NOTE:

Most of the functionality of User Name Mapping has been replaced by Active Directory Lookup. Active Directory Lookup enables Client for NFS and Server for NFS to obtain user identifier (UID) and group identifier (GID) information directly from Active Directory. For information about storing UNIX user data in Active Directory, see documentation for Identity Management for UNIX. For information about enabling Active Directory Lookup, see "Specifying how Server for NFS obtains Windows user and group information" available in the online help.

User Name Mapping Administration

The User Name Mapping administration online help contains information for the following topics:

- Understanding the User Name Mapping component
- Starting and stopping User Name Mapping
- Configuring User Name Mapping
- Securing access to the User Name Mapping server
- Managing maps
- Managing groups

Best practices for User Name Mapping

- Install User Name Mapping on a domain controller.
- Create a User Name Mapping server pool.
- Configure User Name Mapping on a server cluster.
- Make sure User Name Mapping can download users from all domains.
- Refresh data whenever a user is added or changed.
- Place password and group files on the User Name Mapping server.
- Use appropriate permissions to protect password and group files.
- Ensure consistency of group mapping.
- Specify the computers that can access User Name Mapping.

For further details, see the online help for Microsoft Services for Network File System.

Microsoft Services for NFS troubleshooting

The following information on how to troubleshoot issues with Microsoft Services for NFS is available using the online help:

- General issues
- Troubleshooting Server for NFS
- Troubleshooting User Name Mapping

For further details, see the online help for Microsoft Services for Network File System.

Microsoft Services for NFS command-line tools

Table 5 provides a listing of Windows command-line administration tools.

Table 5 MSNFS command-line administration tools

Command	Function
mapadmin	Adds, lists, deletes, or changes user name mappings
mount	Mounts NFS network exports (shares)
nfsadmin	Manages Server for NFS and Client for NFS
nfsshare	Displays, adds, and removes exported NFS shares
nfsstat	Views statistics by NFS operation type
showmount –a	Views users who are connected and what the user currently has mounted
showmount –e	Views exports from the server and their export permissions
unmount	Removes NFS-mounted drives

For further details, see the online help for Microsoft Services for Network File System.

Optimizing Server for NFS performance

The following sources provide useful information on how to optimize performance for Microsoft Services for NFS.

The MSNFS online help covers the following topic areas:

- Adding performance counters
- Monitoring and tuning performance
- Changing the directory cache memory setting

For further details, see the online help for Microsoft Services for Network File System.

A technical paper titled *Performance Tuning Guidelines for Microsoft Services for Network File System* is available at http://www.microsoft.com/technet/interopmigration/unix/sfu/perfnfs.mspx.

Print services for UNIX

Network clients with UNIX-based operating systems that use the client program line printer remote (LPR) can send printing jobs to the line printer daemon (LPD) on the storage server. LPR clients must

comply with Request for Comments (RFC) 1179. The combination of the LPR and LPD are included in print services for UNIX. Print services for UNIX is not preinstalled on the print server.

To install print services for UNIX:

1. Log on as administrator or as a member of the Administrators group.
2. Select **Start** > **Control Panel**, and then click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In the Components list, click **Other Network File and Print Services** (but do not select or clear the check box), and then click **Details**.
5. In the Subcomponents of Other Network File and Print Services list, select **Print Services for UNIX**, if appropriate to the print services that you want to install:

   Print Services for UNIX: This option permits UNIX clients to print to any printer that is available to the print server.

NOTE:

When installing Print Services for UNIX, this automatically installs the LPR port and the TCP/IP Print Server service.

6. Click **OK**, and then click **Next**.
7. Click **Finish**.

Point and print from UNIX to Windows Server 2003

Point-and-Print behavior from UNIX clients to Windows Server 2003 and Windows Storage Server 2003 is similar to the behavior for Windows 95, Windows 98, and Windows Millennium Edition clients, because all these clients create SMB connections. However, the non–Windows operating systems maintain their own driver model, so these clients do not automatically get the driver during Point and Print; they must install the driver locally. Like the Windows 95, Windows 98, and Windows Millennium clients, these non–Windows clients do not receive driver updates from the print server after a driver is initially downloaded. The same connection methods are available: drag and drop, the Add Printer Wizard, referencing a UNC path, or double-clicking the shared printer icon.

Additional resources

Consult the following resources for more information about using and configuring Print Services for UNIX:

- How To: Install and Configure Print Services for UNIX
  http://support.microsoft.com/kb/324078
- How To: Install Print Services for UNIX in Windows Server 2003
  http://support.microsoft.com/?scid=kb;en-us;323421

5 Other network file and print services

This chapter discusses file and print services for NetWare and Macintosh.

File and Print Services for NetWare (FPNW)

File and Print Services for NetWare (FPNW) is one part of the Microsoft software package called Services for NetWare. The most common use of the NetWare network operating system is as a file and print server. FPNW eases the addition of the storage server into a mixed infrastructure by providing a NetWare user interface (UI) to a Windows Storage Server 2003-based server; administrators and users see their same, familiar NetWare UI. Additionally, the same logon for clients is maintained without a need for any client configuration changes.

This service also provides the ability to create Novell volumes, which are actually NTFS shares, from which users can map drives and access resources. Novell Login scripts are supported on the storage server or through an existing NDS (Novell Directory Services) account. This requires no changes or additions to the software on the NetWare client computers.

NOTE:

FPNW is not a clusterable protocol. With FPNW on both nodes of a cluster, the shares do not fail over because the protocol is not cluster-aware.

NOTE:

IPX/SPX protocol is required on the Novell servers.

Installing Services for NetWare

The installation of FPNW on the storage server allows for a smooth integration with existing Novell servers. FPNW allows a Windows Storage Server 2003 based server to emulate a NetWare file and print server to users, clients, and administrators. This emulation allows authentication from Novell clients, the use of Novell logon scripts, the creation of Novell volumes (shares), the use of Novell file attributes, and many other Novell features.

Information on Microsoft Directory Synchronization Services and the File Migration Utility can be found at

http://www.microsoft.com/WINDOWS2003/guide/server/solutions/NetWare.asp

To install Services for NetWare, , naviagte to the `c:\hpnas\components\SFN5.003SP2` folder and run the `FPNW 5.02.exe` setup executable file.

Managing File and Print Services for NetWare

FPNW resources are managed through Server Manager. Server Manager can be used to modify FPNW properties and manager shared volumes.

Use File and Print Services for NetWare to:

- Access files, modify file settings and permissions from Computer Management, and use third party tools that can be used with NetWare servers.
- Create and manage user accounts by using Active Directory Users and Computers.
- Perform secured log-ons.
- Support packet burst and Large Internet Packet (LIP).
- Support NetWare locking and synchronization primitives that are used by some NetWare-specific applications.
- Support long file names, compatible with OS/2 long file name (LFN) support.

File and Print Services for NetWare does not support the following NetWare groups and functions:

- Workgroup Managers
- Accounting
- User disk volume restrictions
- Setting Inherited Rights Masks (IRMs)
- NetWare loadable modules
- Transaction Tracking System (TTS)

To access FPNW:

1. From the desktop of the storage server, select **Start** > **Settings** > **Control Panel > Administrative Tools > Server Manager**.

2. Select **FPNW**, and then click **Properties**.

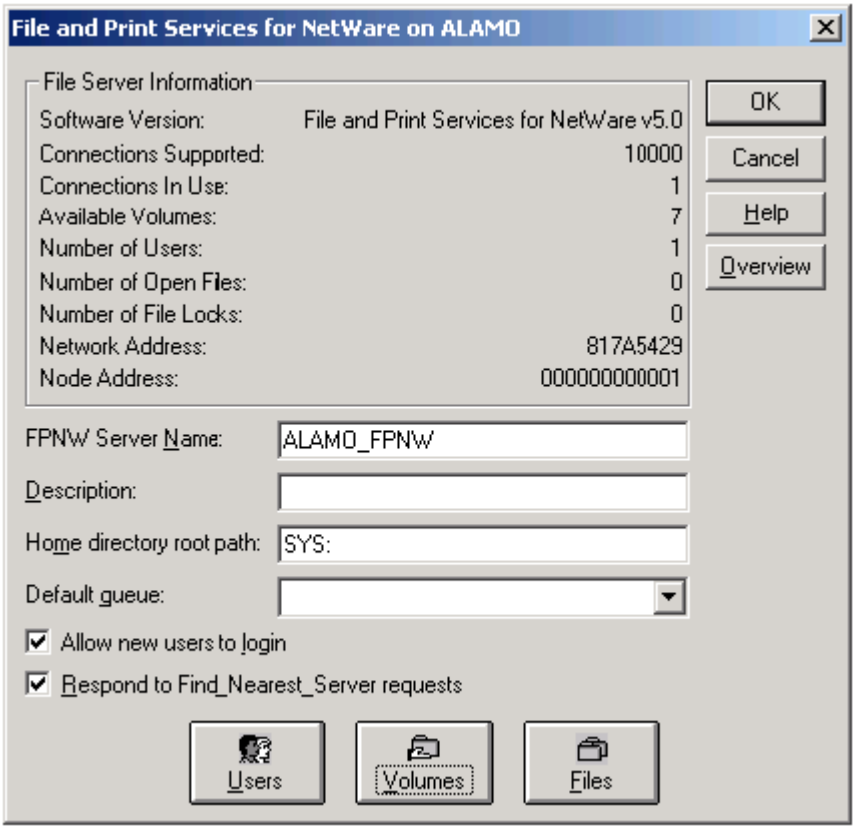

Figure 19 File and Print Services for NetWare dialog box

3. Enter an FPNW Server Name and Description.

   This server name must be different from the server name used by Windows or LAN Manager-based clients. If changing an existing name, the new name is not effective until stopping and restarting FPNW. For example, in Figure 19 the Windows server name is Alamo and the FPNW server name is Alamo_FPNW.

4. Indicate a Home directory root path.

   This path is relative to where the Sysvol volume is installed. This is the root location for the individual home directories. If the directory specified does not already exist, it must first be created.

5. Click **Users** to:

   See connected users, disconnect users, send broadcast messages to all users connected to the server, and to send a message to a specific user.

6. Click **Volumes** to:

   See users connected to specific volume and to disconnect users from a specific volume.

7. Click **Files** to:

   View open files and close open files.

Creating and managing NetWare users

To use Services for NetWare, the Novell clients must be entered as local users on the storage server.

Adding local NetWare users

1. From the storage server desktop, click the **Management Console** icon, click **Core Operating System**, and then click **Local Users and Groups**.
2. Right-click the **Users** folder, and then click **New User**.

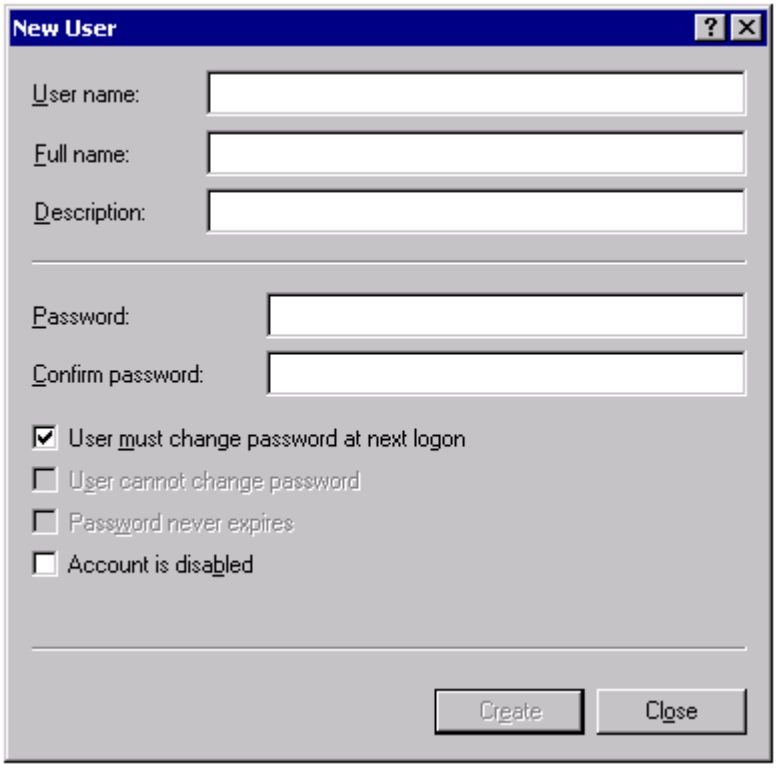

Figure 20 New User dialog box

3. Enter the user information, including the user's User name, Full name, Description, and Password.
4. Click **Create**.
5. Repeat these steps until all NetWare users have been entered.

Enabling local NetWare user accounts

1. In the **Users** folder (MC, Core Operating System, Local Users and Groups), right-click an NCP client listed in the right pane of the screen, and then click **Properties**.

2. Click the **NetWare Services** tab.

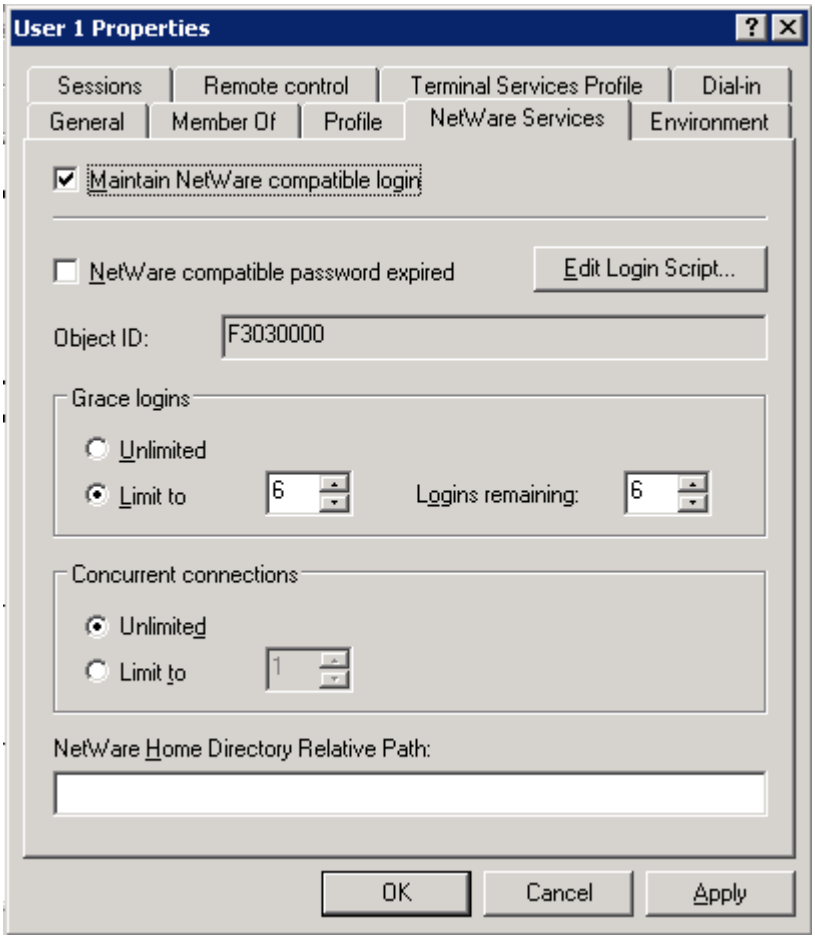

Figure 21 NetWare Services tab

3. Select **Maintain NetWare compatible login**.
4. Set other NetWare options for the user, and then click **OK**.

NOTE:

The installation of File and Print Services for NetWare also creates a supervisor account, which is used to manage FPNW. The supervisor account is required if the storage server was added as a bindery object into NDS.

Managing NCP volumes (shares)

NCP file shares are created the same way as other file shares; however, there are some unique settings. NCP shares can be created and managed using Server Manager.

NOTE:

NCP shares can be created only after FPNW is installed. See the previous section "Installing Services for Netware" for instructions on installing FPNW.

Creating a new NCP share

To create a new file share:

1. From the storage server desktop, select **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Server Manager**.
2. Select **File and Print Service for NetWare**> **Shared Volumes**.
3. Click **Create Volume**.
4. Specify the volume name and path.
5. Click **Permissions** to set permissions.
6. Click **Add** to add additional users and groups, and to set their permissions.
7. Highlight the desired user or group, and then click **Add**.
8. Select the Type of Access in the drop down list.

   Type of Access can also be set from the Access Through Share Permissions dialog box.
9. Click **OK** when all users and groups have been added.
10. Click **OK** in the **Create Volume** dialog box.
11. Click **Close**.

Modifying NCP share properties

To modify a file share:

1. From the storage server desktop, select **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Server Manager**.
2. Select **File and Print Services for NetWare** > **Shared Volumes**.
3. Highlight the volume to modify.
4. Click **Properties**.

Print Services for NetWare

With File and Print Services for NetWare installed, the print server appears to a NetWare client as a NetWare 3.x–compatible print server. Print services presents the same dialog boxes to the client as a NetWare-based server uses to process a print job from a client. A user can display and search for printers on the print server just like in a NetWare environment.

Installing Print Services for NetWare

See the previous section "Installing Services for Netware" for information on installing Print Services for NetWare.

Point and Print from Novell to Windows Server 2003

Point-and-Print behavior from Novell clients to Windows Server 2003 and Windows Storage Server 2003 is similar to the behavior for Windows 95, Windows 98, and Windows Millennium Edition clients, because all these clients create SMB connections. However, the non–Windows operating systems maintain their own driver model, so these clients do not automatically get the driver during Point and Print—they must install the driver locally. Like the Windows 95, Windows 98, and Windows

Millennium clients, these non–Windows clients do not receive driver updates from the print server after a driver is initially downloaded. The same connection methods are available: drag and drop, the Add Printer Wizard, referencing a UNC path, or double-clicking the shared printer icon.

Additional resources

For more information about using and configuring File and Print Services for NetWare, see the online help.

AppleTalk and file services for Macintosh

The AppleTalk network integration allows the storage server to share files and printers between your server and any Apple Macintosh clients that are connected to your network. After installing Microsoft Windows Services for Macintosh, the administrator can use the AppleTalk protocol to configure the storage server to act as an AppleTalk server. The AppleTalk protocol is the communications protocol used by clients running a Macintosh operating system. The Macintosh computers need only the Macintosh OS software to function as clients; no additional software is required.

AppleTalk network integration simplifies administration by maintaining just one set of user accounts instead of separate user accounts, for example, one on the Macintosh server and another on the computer running Windows server software.

Installing the AppleTalk protocol

1. From the desktop of the storage server, select **Start > Settings > Network Connections**. Right-click **Local Area Connection**, and then click **Properties**.
2. Click **Install**.
3. Select **Protocol**, and then click **Add**.
4. Select **AppleTalk Protocol**, and then click **OK**.

Installing File Services for Macintosh

To install File Services for Macintosh, perform the following steps:

1. Access the desktop on the storage server.
2. Open **Add or Remove Programs** from the Control Panel.
3. Click **Add or Remove Windows Components**.
4. Double-click **Other Network File and Print Services**.
5. Select **File Services for Macintosh**, and then click **OK**.
6. Click **Next**.
7. Click **Finish**.

Completing setup of AppleTalk protocol and shares

See the online help to complete the following setup and configurations tasks:

- To set up AppleTalk protocol properties

  AppleTalk shares can be set up only after AppleTalk Protocol and File Services for Macintosh have been installed on the storage server.

△ CAUTION:

AppleTalk shares should not be created on clustered resources because data loss can occur due to local memory use.

- To set up AppleTalk shares
- To configure AppleTalk sharing properties
- To allow client permission to an AppleTalk share

  If AppleTalk is enabled for your server configuration, specify which AppleTalk clients are granted access to each share. Access can be granted or denied on the basis of client host name. Access can also be granted or denied on the basis of client groups, where a client group contains one or more client host names.

Print services for Macintosh

Macintosh clients can send print jobs to a print server when Print Server for Macintosh is installed on the server. To the Macintosh-based client, the print server or FPA appears to be an AppleTalk printer on the network, and no reconfiguration of the client is necessary.

Installing Print Services for Macintosh

Consult the following resource for information about installing Print Services for Macintosh:

- How To: Install Print Services for Macintosh in Windows Server 2003
  http://support.microsoft.com/?scid=kb;en-us;323421

Point and Print from Macintosh to Windows Server 2003

Point-and-Print behavior from Macintosh clients to Windows Server 2003 or Windows Storage Server 2003 is similar to the behavior for Windows 95, Windows 98, and Windows Millennium Edition clients, because all these clients create SMB connections. However, the non-Windows operating systems maintain their own driver model, so these clients do not automatically get the driver during Point and Print; they must install the driver locally. Like the Windows 95, Windows 98, and Windows Millennium clients, these non-Windows clients do not receive driver updates from the print server after a driver is initially downloaded. The same connection methods are available: drag and drop, the Add Printer Wizard, referencing a UNC path, or double-clicking the shared printer icon.

6 Enterprise storage servers

Some HP ProLiant Storage Servers use the Microsoft® Windows® Unified Data Storage Server 2003 operating system. This operating system provides unified storage server management capabilities, simplified setup and management of storage and shared folders, and support for Microsoft iSCSI Software Target. It is specially tuned to provide optimal performance for network-attached storage and provides significant enhancements in share and storage management scenarios, as well as integration of storage server management components and functionality. This chapter describes features of the Microsoft® Windows® Unified Data Storage Server 2003 operating system.

NOTE:

Not all HP ProLiant Storage Servers use the Microsoft® Windows® Unified Data Storage Server 2003, Enterprise x64 Edition operating system. See the HP ProLiant Storage Server QuickSpecs to determine if your storage server runs this operating system.

IMPORTANT:

The Microsoft® Windows® Unified Data Storage Server 2003, Enterprise x64 Edition operating system is designed to support 32–bit applications without modification; however, any 32–bit applications that are run on this operating system should be thoroughly tested before releasing the storage server to a production environment.

Windows Server Remote Administration Applet

Remote administration from non-Microsoft computers uses the Windows Server Remote Administration Applet and is accessed from a browser. The browser on the client computer can be any of the following:

- Firefox version 1.0.6 (or later)
- Mozilla version 1.7.11 (or later)

Use of Windows Server Remote Administration Applet is supported by clients running Java 2 Runtime Environment, version 1.4.2 on any of the following:

- A computer running a Windows operating system and Internet Explorer 6 or later browser
- Any of the following non-Microsoft operating systems: Red Hat Enterprise Linux 3 WS, Red Hat Enterprise Linux 4 WS, SuSE Linux Enterprise Server 9, SuSE Linux Enterprise Server 10

Establishing a connection is done directly through the browser. Windows Server Remote Administration Applet does not support sound redirection, printer or port redirection, or automatically starting applications.

To establish a browser-based connection to Windows Unified Data Storage Server 2003

1. Open the browser on the client computer.

2. Type the network name or the network IP address of the storage server followed by `/admin` (for example, `http://myStorageServer/admin`).
3. In **Remote Administration Desktop**, provide the appropriate credentials.

NOTE:

Administrative credentials are not required to establish a browser-based connection, but are required to manage the storage server.

NOTE:

If Java Runtime Environment (JRE) is not installed correctly, the Additional plug-ins are required to display all the media on this page message may be displayed. For information about installing JRE on a non-Microsoft system, see Installation Instructions (http://go.microsoft.com/fwlink/?LinkId=70026).

Microsoft iSCSI Software Target

The Microsoft iSCSI Software Target snap-in is a standard feature of Windows Unified Data Storage Server 2003. This snap-in makes it possible not only for the storage server to connect to remote iSCSI targets, but also to serve as an iSCSI target. With Microsoft iSCSI Software Target, you can create and manage iSCSI targets, create and manage disks for storage, and implement backup and recovery support using snapshots.

Virtual disk storage

The disks you create using iSCSI Software Target are iSCSI virtual disks, which are files in the virtual hard disk (VHD) format. These virtual disks offer flexible and effective storage. They are dynamically extendable to provide extra capacity on demand, enable efficient storage utilization, and minimize the time required to create new disks and the down time typically required to install new disks.

Snapshots

To facilitate backup and recovery operations, you can schedule and create snapshots of iSCSI virtual disks. A snapshot is a point-in-time, read-only copy of an iSCSI virtual disk. Snapshots are typically used as interim copies of data that has been modified since the most recent backup. Snapshots offer the following advantages:

- Snapshots can be scheduled to be created automatically.
- Snapshots are space-efficient because they are differential copies.
- It is not necessary to close files or stop programs when creating snapshots, so application servers can continue servicing clients without disruption.
- Each snapshot is typically created in less than one minute—regardless of the amount of data.
- Snapshots are useful for fast system recovery of files and volumes, in case of accidental data deletion by a user, overwritten data, or data corruption resulting from a malicious program.
- Snapshots can be mounted locally or exported to facilitate backup and recovery operations.

Wizards

To support creation and management of iSCSI targets, virtual disks, and snapshots, the iSCSI Software Target snap-in provides several wizards.

Create iSCSI Target Wizard

This section describes how to create an iSCSI Target using the Create iSCSI Target Wizard.

1. Log on to the storage server using an account with administrative privileges.
2. Open the Microsoft iSCSI Software Target MMC snap-in by clicking **Start > Programs > Administrative Tools > Microsoft iSCSI Software Target**.
3. Click the **iSCSI Targets** node. On the details view (right pane), right-click and select **Create iSCSI Target**.
4. Click **Next** on the **Welcome** page of the wizard.
5. On the **iSCSI Target Identification** page, type a name and description for the iSCSI Target and then click **Next**.
6. On the **iSCSI Initiators Identifiers** page, type the iSCSI Qualified Name (IQN) of the iSCSI initiator requesting access to the iSCSI Target in the **IQN identifier** field. The IQN is found on the **General** tab of the Microsoft iSCSI Initiator interface.

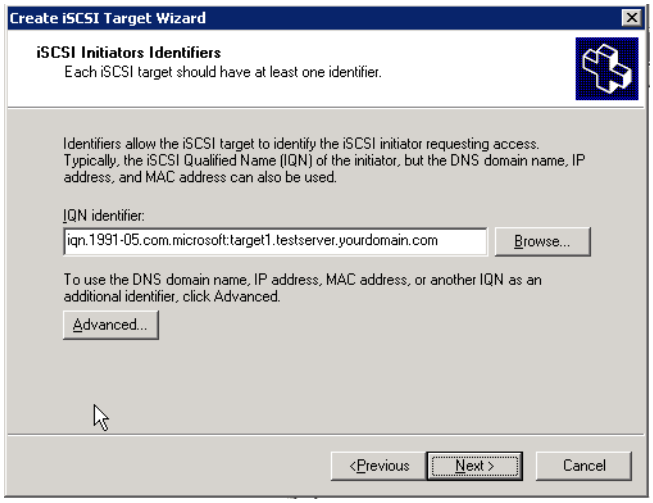

Figure 22 iSCSI Initiators Identifiers page

7. To enter additional identifiers, or if you are using an identifier other than an IQN (DNS domain name, IP address, or MAC address):
   a. Click **Advanced**.
   b. On the **Advanced Identifiers** page, click **Add**.
   c. Select the identifier type from the **Identifier Type** list and type the identifier in the **Value** field.
   d. Repeat steps b and c for each identifier you want to add.
   e. Click **OK**.
   f. Click **OK** again to close the **Advanced Identifiers** page.

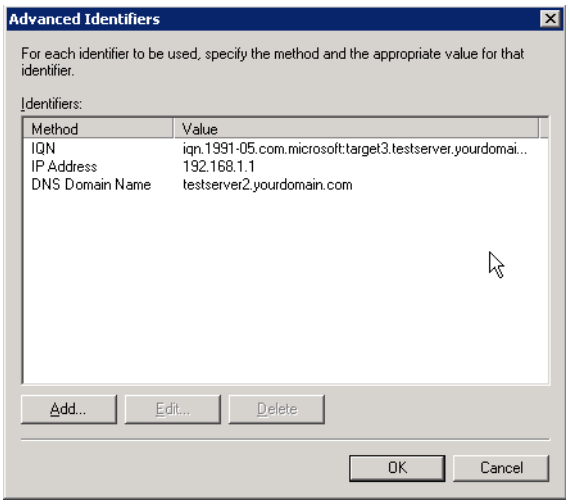

Figure 23 Advanced Identifiers page

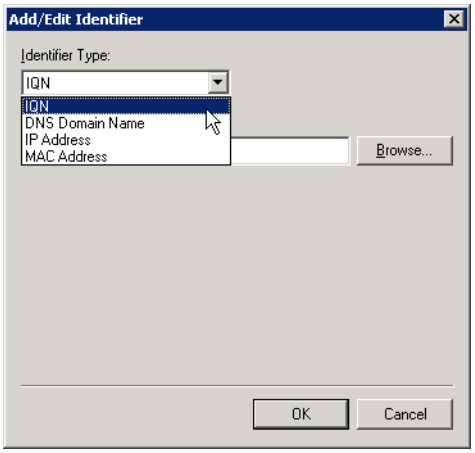

Figure 24 Add/Edit Identifier page

8. Click **Next**.
9. Click **Finish** to complete the wizard and create the iSCSI Target.

Create Virtual Disk Wizard

This section describes how to create an iSCSI Virtual Disk using the Create Virtual Disk Wizard.

NOTE:

In order to create iSCSI Virtual Disks, it is required that physical disks are formatted as NTFS.

1. In the **Microsoft iSCSI Software Target** MMC snap-in, click the **Devices** node.
2. On the details view (right pane) of the **Devices** node, right-click a volume and select **Create Virtual Disk**.
3. Click **Next** on the **Welcome** page of the wizard.
4. On the **File** page, specify the full path to use as the virtual disk and click **Next**.
5. On the **Size** page, specify the size to use for the virtual disk and click **Next**. If the file already exists, you cannot specify a new size.

6. Enter a description for the iSCSI virtual disk (optional) and click **Next**.
7. On the **Access** page, click **Add** to assign the iSCSI virtual disk to an iSCSI Target.
8. On the **Add Targets** dialog box, select a Target and click **OK**.
9. Click **Finish** to complete the wizard and create the iSCSI virtual disk.

NOTE:

If you delete a virtual disk, it is removed from the iSCSI Software Target MMC snap-in, but the virtual disk file (.vhd) is not removed from the physical disk. In order to permanently remove the virtual disk file, locate the file on the physical disk using Windows Explorer and manually delete it.

Import Virtual Disk Wizard

This section describes how to import a virtual disk using the Import Virtual Disk Wizard.

1. In the **Microsoft iSCSI Software Target** MMC snap-in, click the **Devices** node.
2. On the details view (right pane) of the **Devices** node, right-click a volume and select **Import Virtual Disk**.
3. Click **Next** on the **Welcome** page of the wizard.
4. On the **Files** page, click **Browse**, navigate to the virtual disk file (.vhd) you want to import, select it, and then click **OK**.
5. Repeat step 4 for each virtual disk you want to import.
6. Click **Next** and then click **Finish** to complete the wizard and import the virtual disk(s).

Extend Virtual Disk Wizard

This section describes how to extend a virtual disk using the Extend Virtual Disk Wizard.

1. In the **Microsoft iSCSI Software Target** MMC snap-in, click the **Devices** node.
2. On the details view (right pane) of the **Devices** node, right-click a virtual disk and select **Extend Virtual Disk**.
3. Click **Next** on the **Welcome** page of the wizard.
4. On the **Size** page, type the amount of space you want to add to the virtual disk in the **Additional virtual space capacity** field and then click **Next**.
5. Click **Finish** to complete the wizard and extend the virtual disk.

Schedule Snapshot Wizard

This section describes how to schedule a snapshot using the Schedule Snapshot Wizard.

1. In the **Microsoft iSCSI Software Target** MMC snap-in, expand the **Snapshots** node.
2. Right-click **Schedule** and select **Create Schedule**.
3. Click **Next** on the **Welcome** page of the wizard.
4. On the **Schedule Actions** page, specify whether the snapshots should be mounted locally or not.
5. On the **Name** page, type a name for the snapshot and then click **Next**.
6. On the **Virtual Disks** page, specify the virtual disks to include in the snapshot schedule.
7. On the **Frequency** page, select how often snapshots should be taken.
8. On the **Schedule** page, specify snapshot details according to the frequency selected on the previous page and then click **Next**.

9. Click **Finish** to complete the wizard and schedule snapshots.

Hardware provider

To support advanced management of iSCSI virtual disks and snapshots, you can use the Microsoft iSCSI Software Target Virtual Disk Service Hardware Provider, which comes preinstalled on the HP ProLiant Storage Server.

Microsoft Windows Server 2003 introduced Virtual Disk Service (VDS), a set of application programming interfaces (APIs) that provides a single interface for managing disks. VDS provides an end-to-end solution for managing storage hardware and disks, and for creating volumes on those disks. The Microsoft iSCSI Software Target VDS Hardware Provider is required to manage virtual disks on a storage subsystem.

You install the Microsoft iSCSI Software Target VDS Hardware Provider on each iSCSI initiator computer running a storage management application (such as Storage Manager for SANs) that uses the hardware provider to manage storage.

- Microsoft iSCSI Software Target Volume Shadow Copy Service Hardware Provider

iSCSI snapshots are created using Volume Shadow Copy Service and a storage array with a hardware provider designed for use with Volume Shadow Copy Service. A Microsoft iSCSI Software Target VSS Hardware Provider is required to create transportable snapshots of iSCSI virtual disks and create application consistent snapshots from iSCSI initiators.

You install this hardware provider on the iSCSI initiator server and the server that is to perform backups. The backup software you use must support transporting snapshots.

Cluster support

In a cluster with servers running Windows Unified Data Storage Server 2003, Enterprise Edition and using an external storage array as the shared cluster disk, you can use iSCSI Software Target to share highly available storage. To do this, use Cluster Administrator to configure the iSCSI target as a Generic Service cluster resource. iSCSI virtual disks can then be created from the generic cluster disk and exported to iSCSI initiators.

IMPORTANT:

A single-server iSCSI software target cluster configuration does not provide the redundant components of a hardware-based shared disk resource, making it a potential single point of failure. In most cases, this type of configuration does not provide the level of availability typically required in a production environment.

For detailed instructions on how to set up a cluster using Microsoft iSCSI Target as the shared-cluster disk provider, see the HP white paper *Using Microsoft iSCSI Software Target to Provide Shared-Disk Resources to Clusters* at http://h71028.www7.hp.com/ERC/downloads/4AA1-0720ENW.pdf.

For detailed instructions on how to set up an iSCSI software target cluster, see the HP white paper *Configuring Microsoft iSCSI Software Target in a Microsoft Cluster* at http://h71028.www7.hp.com/ERC/downloads/4AA1-2898ENW.pdf.

7 Cluster administration

NOTE:

Not all HP ProLiant Storage Servers can be clustered. See the HP ProLiant Storage Server QuickSpecs to determine if your storage server can be clustered. Windows Storage Server 2003 Release 2 clusters can include up to eight nodes.

One important feature of the HP ProLiant Storage Server clusterable models is that they can operate as a single node or as a cluster. This chapter discusses cluster installation and cluster management issues.

Cluster overview

Up to eight server nodes can be connected to each other and deployed as a no single point of failure (NSPOF) cluster. Utilizing a private network allows communication amongst themselves in order to track the state of each cluster node. Each node sends out periodic messages to the other nodes; these messages are called heartbeats. If a node stops sending heartbeats, the cluster service fails over any resources that the node owns to another node. For example, if the node that owns the Quorum disk is shut down for any reason, its heartbeat stops. The other nodes detect the lack of the heartbeat and another node takes over ownership of the Quorum disk and the cluster.

Clustering servers greatly enhances the availability of file serving by enabling file shares to fail over to additional storage servers if problems arise. Clients see only a brief interruption of service as the file share resource transitions from one server node to the other.

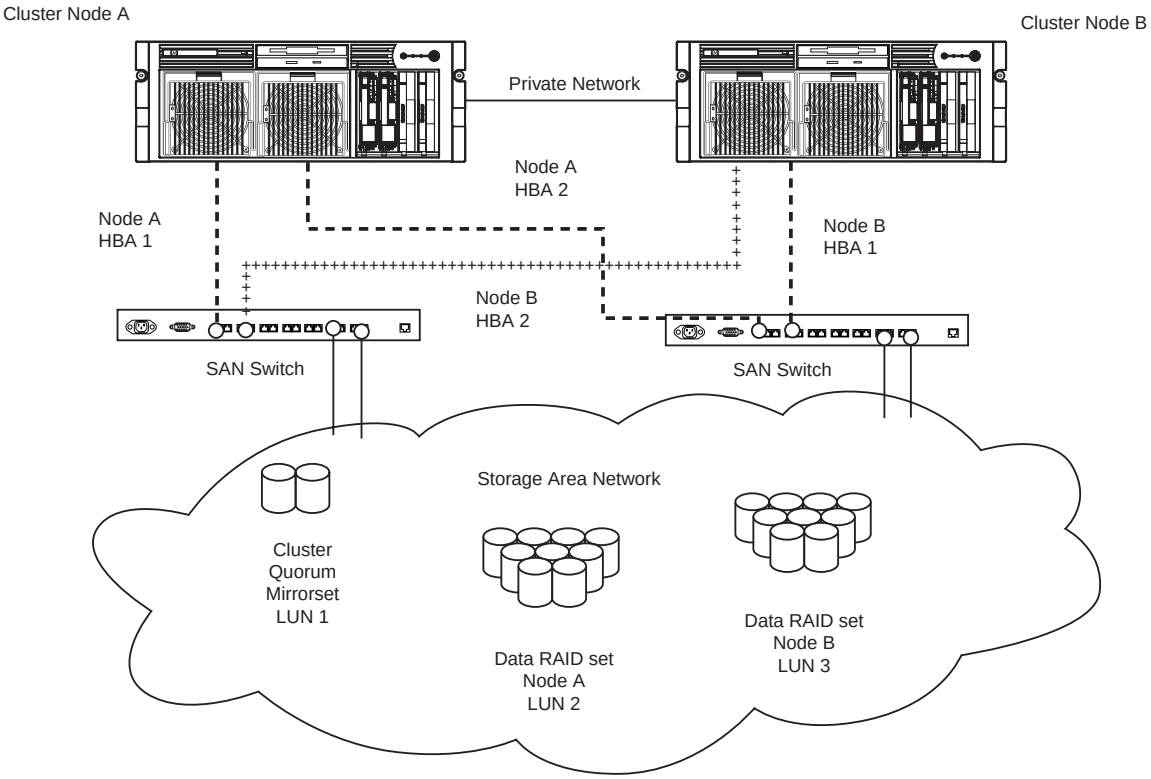

Figure 25 Storage server cluster diagram

Cluster terms and components

Nodes

The most basic parts of a cluster are the servers, referred to as nodes. A server node is any individual server in a cluster, or a member of the cluster.

Resources

Hardware and software components that are managed by the cluster service are called cluster resources. Cluster resources have three defining characteristics:

- They can be brought online and taken offline.
- They can be managed in a cluster.
- They can be owned by only one node at a time.

Examples of cluster resources are IP addresses, network names, physical disk resources, and file shares. Resources represent individual system components. These resources are organized into groups and managed as a group. Some resources are created automatically by the system and other resources must be set up manually. Resource types include:

- IP address resource
- Cluster name resource
- Cluster quorum disk resource
- Physical disk resource
- Virtual server name resources

- CIFS file share resources
- NFS file share resources
- FTP file share resources
- iSCSI resources

Cluster groups

Cluster resources are placed together in cluster groups. Groups are the basic unit of failover between nodes. Resources do not fail over individually; they fail over with the group in which they are contained.

Virtual servers

A virtual server is a cluster group that consists of a static IP Address resource and a Network Name resource. Several virtual servers can be created. By assigning ownership of the virtual servers to the different server nodes, the processing load on the storage servers can be distributed between the nodes of a cluster.

The creation of a virtual server allows resources dependent on the virtual server to fail over and fail back between the cluster nodes. Cluster resources are assigned to the virtual server to ensure non-disruptive service of the resources to the clients.

Failover and failback

Failover of cluster groups and resources happens:

- When a node hosting the group becomes inactive.
- When all of the resources within the group are dependent on one resource, and that resource fails.
- When an administrator forces a failover.

A resource and all of its dependencies must be located in the same group so that if a resource fails over, all of its dependent resources fail over.

When a resource is failed over, the cluster service performs certain procedures. First, all of the resources are taken offline in an order defined by the resource dependencies. Secondly, the cluster service attempts to transfer the group to the next node on the preferred owner's list. If the transfer is successful, the resources are brought online in accordance with the resource dependency structure.

The system failover policy defines how the cluster detects and responds to the failure of individual resources in the group. After a failover occurs and the cluster is brought back to its original state, failback can occur automatically based on the policy. After a previously failed node comes online, the cluster service can fail back the groups to the original host. The failback policy must be set before the failover occurs so that failback works as intended.

Quorum disk

Each cluster must have a shared disk called the Quorum disk. The Quorum disk is the shared storage used by the cluster nodes to coordinate the internal cluster state. This physical disk in the common cluster disk array plays a critical role in cluster operations. The Quorum disk offers a means of persistent storage. The disk must provide physical storage that can be accessed by all nodes in the cluster. If a node has control of the quorum resource upon startup, it can initiate the cluster. In addition, if the node can communicate with the node that owns the quorum resource, it can join or remain in the cluster.

The Quorum disk maintains data integrity by:

- Storing the most current version of the cluster database
- Guaranteeing that only one set of active communicating nodes is allowed to operate as a cluster

Cluster concepts

Figure 26 illustrates a typical cluster configuration with the corresponding storage elements. The diagram progresses from the physical disks to the file shares, showing the relationship between both the cluster elements and the physical devices underlying them. While the diagram only illustrates two nodes, the same concepts apply for multi-node deployments.

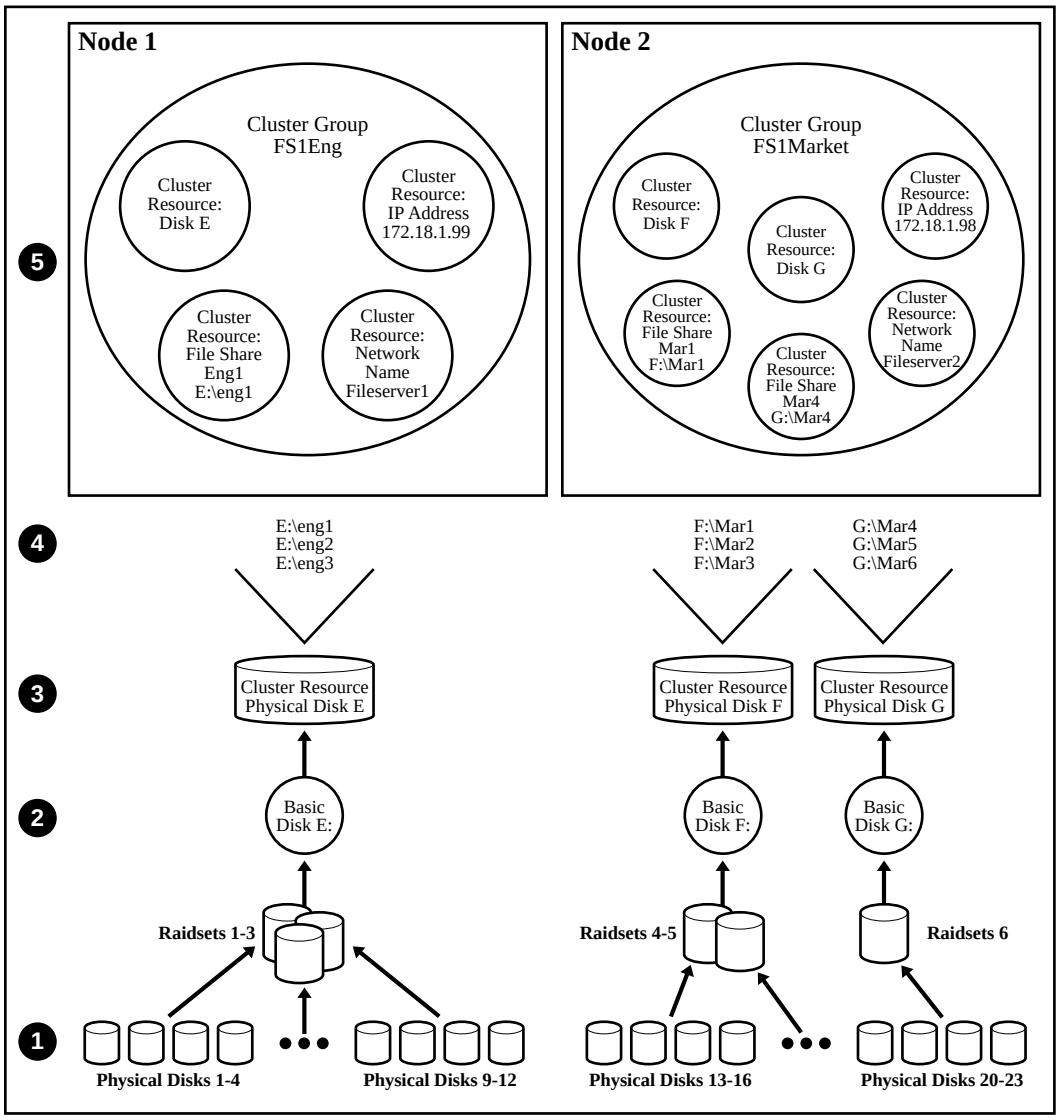

Figure 26 Cluster concepts diagram

Sequence of events for cluster resources

The sequence of events in the diagram includes:

1. Physical disks are combined into RAID arrays and LUNs.
2. LUNS are designated as basic disks, formatted, and assigned a drive letter via Disk Manager.

3. Physical Disk resources are created for each basic disk inside Cluster Administrator.
4. Directories and folders are created on assigned drives.
5. Cluster components (virtual servers, file shares) are created, organized in groups, and placed within the folders using Cluster Administrator exclusively.

Hierarchy of cluster resource components

Figure 26 depicts the cluster resource hierarchy as follows:

- Physical Disk resources are placed in a cluster group and relate to the basic disk. When a Physical Disk resource is created through Cluster Administrator, the resource should be inserted into an existing cluster group or a corresponding group should be created for the resource to reside in.
- File share resources are placed in a group and relate to the actual directory on the drive on which the share is being created.
- An IP Address resource is formed in the group and relates to the IP address by which the group's virtual server is identified on the network.
- A Network Name resource is formed in the group and relates to the name published on the network by which the group is identified.
- The Group is owned by one of the nodes of the cluster, but may transition to the other nodes during failover conditions.

The diagram illustrates a cluster containing two nodes. Each node has ownership of one group. Contained within each group are file shares that are known on the network by the associated Network Name and IP address. In the specific case of Node1, file share Eng1 relates to `E:\Eng1`. This file share is known on the network as `\\Fileserver1\Eng1` with an IP address of `172.18.1.99`.

For cluster resources to function properly, two very important requirements should be adhered to:

- Dependencies between resources of a group must be established. Dependencies determine the order of startup when a group comes online. In the above case, the following order should be maintained:
  1. File Share—Dependent on Physical Disk Resource and Network Name
  2. Network Name—Dependent on IP Address

  Failure to indicate the dependencies of a resource properly may result in the file share attempting to come online prior to the physical disk resource being available, resulting in a failed file share.
- Groups should have a Network Name resource and an IP Address resource. These resources are used by the network to give each group a virtual name. Without this virtual reference to the group, the only way to address a share that is created as a clustered resource is by node name. Physical node names do not transition during a failover, whereas virtual names do.

For example, if a client maps a network share to `\\Node1\Eng1` instead of `\\Fileserver1\Eng1`, when Node1 fails and Node2 assumes ownership, the map will become invalid because the reference in the map is to `\\Node1`. If the map were created to the virtual name and Node1 were to fail, the map would still exist when the group associated with Eng1 failed over to Node2.

The previous diagram is an example and is not intended to imply limitations of a single group or node. Groups can contain multiple physical disks resources and file shares and nodes can have multiple groups, as shown by the group owned by Node2.

Cluster planning

Requirements for taking advantage of clustering include:

- Storage planning
- Network planning
- Protocol planning

Storage planning

For clustering, a basic disk must be designated for the cluster and configured as the Quorum disk.

Additional basic disks are presented to each cluster node for data storage as physical disk resources. The physical disk resources are required for the basic disks to successfully work in a cluster environment, protecting it from simultaneous access from each node.

The basic disk must be added as a physical disk resource to an existing cluster group or a new cluster group needs to be created for the resource. Cluster groups can contain more than one physical disk resource depending on the site-specific requirements.

NOTE:

The LUN underlying the basic disk should be presented to only one node of the cluster using selective storage presentation or SAN zoning, or having only one node online at all times until the physical resource for the basic disk is established.

In preparing for the cluster installation:

- All software components listed in the *HP ProLiant Storage Server SAN Connection and Management* white paper (located on the HP web site at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00663737/c00663737.pdf) must be installed and the fiber cables attached to the HBAs before the cluster installation is started.
- All shared disks, including the Quorum disk, must be accessible from all nodes. When testing connectivity between the nodes and the LUN, only one node should be given access to the LUN at a time.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Network planning

Clusters require more sophisticated networking arrangements than a stand alone storage server. A Windows NT domain or Active Directory domain must be in place to contain the cluster names, virtual server names, and user and group information. A cluster cannot be deployed into a non domain environment.

All cluster deployments have at least six network addresses and four network names:

- The cluster name (Unique NETBIOS Name) and IP address
- Node A's name and IP address
- Node B's name and IP address
- At least one virtual server name and IP address for virtual server A
- Cluster Interconnect static IP addresses for Node A and Node B

In multi-node deployments, additional network addresses are required. For each additional node, three static IP addresses are required.

Virtual names and addresses are the only identification used by clients on the network. Because the names and addresses are virtual, their ownership can transition from one node to the other during a failover, preserving access to the resources in the cluster group.

A cluster uses at least two network connections on each node:

- The private cluster interconnect or "heartbeat" crossover cable connects to one of the network ports on each cluster node. In more than two node deployments, a private VLAN on a switch or hub is required for the cluster interconnect.
- The public client network subnet connects to the remaining network ports on each cluster node. The cluster node names and virtual server names have IP addresses residing on these subnets.

NOTE:

If the share is to remain available during a failover, each cluster node must be connected to the same network subnet. It is impossible for a cluster node to serve the data to a network to which it is not connected.

Protocol planning

Not all file sharing protocols can take advantage of clustering. If a protocol does not support clustering, it will not have a cluster resource and will not failover with any cluster group. In the case of a failover, a client cannot use the virtual name or virtual IP address to access the share since the protocol cannot failover with the cluster group. The client must wait until the initial node is brought back online to access the share.

HP recommends placing cluster aware and non cluster aware protocols on different file shares.

Table 6 Sharing protocol cluster support

Protocol	Client Variant	Cluster Aware (supports failover)	Supported on cluster nodes
CIFS/SMB	Windows NT Windows 2000 Windows 95 Windows 98 Windows ME	Yes	Yes
NFS	UNIX Linux	Yes	Yes
HTTP	Web	No	Yes
FTP	Many	Yes	Yes
NCP	Novell	No	Yes
AppleTalk	Apple	No	No
iSCSI	Standards-based iSCSI initiator	Yes	Yes

NOTE:

AppleTalk is not supported on clustered disk resources. AppleTalk requires local memory for volume indexing. On failover events, the memory map is lost and data corruption can occur.

Preparing for cluster installation

This section provides the steps necessary to cluster HP ProLiant Storage Servers.

Before beginning installation

Confirm that the following specifications have been met before proceeding:

- The procedures in the *HP ProLiant Storage Server SAN Connection and Management* white paper (located on the HP web site at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00663737/c00663737.pdf) must be completed and all the necessary software components for connecting to the desired storage must be installed before the configuration of cluster services.
- The Quorum disk has been created from shared storage and is at least 50 MB. (500 MB is recommended.) Additional LUNs may also be presented for use as shared disk resources.
- Cluster configurations should be deployed with dual data paths for high availability. Dual data paths from each node enable a path failure to occur that does not force the failover of the node. Clusters can be configured with single path, but if a failure in the path does occur, all of the node resources will be failed to the non-affected node.

Using multipath data paths for high availability

HP recommends that cluster configurations be deployed with dual data paths for high availability. Clusters can be configured with single path, but if a failure in the path occurs, all of the node resources will be failed to the non-affected node. Pathing software is required in configurations where multipathing to the storage is desired or required. Multipathing software allows for datapath failure to occur without forcing a node failover. See the *HP ProLiant Storage Server SAN Connection and Management* white paper (located on the HP web site at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00663737/c00663737.pdf) for additional information on multipath software.

Enabling cluster aware Microsoft Services for NFS (optional)

The server comes with Microsoft Services for Network File System (NFS) preinstalled. Microsoft has identified an issue that requires that NFS be uninstalled prior to installing clustering on the storage server. This only applies if you want NFS share resources on the storage server within the clustering environment. After clustering is installed, then NFS can be installed if you desire this feature.

NOTE:

Failure to uninstall Microsoft NFS prior to the installation of clustering results in no NFS resource types being available in the clustering environment.

To uninstall Microsoft NFS:

1. From the storage server desktop, select **Start > Settings > Control Panel > Add or Remove Programs**. The Add or Remove Programs window is displayed.

2. On the left side of the window, select **Add/Remove Windows Components**. The Windows Components Wizard appears.
3. Select **Other Network File and Print Services** and click the **Details** button. The Other Network File and Print Services window is displayed.
4. Uncheck the **Microsoft Services for NFS** subcomponent.
5. Click **OK**, then **Next**, followed by **Finish**.

NOTE:

Uninstalling Microsoft Services for NFS removes two primary services:

- Server for NFS
- User Name Mapping

After setting up clustering, should you choose to reinstall Microsoft NFS, follow these steps:

1. Select **Start > Settings > Control Panel > Add or Remove Programs**. The Add or Remove Programs window is displayed.
2. On the left side of the window, select **Add/Remove Windows Components**. The Windows Components Wizard appears.
3. Select **Other Network File and Print Services** and click the **Details** button. The Other Network File and Print Services window is displayed.
4. Check the **Microsoft Services for NFS** subcomponent.
5. Click **OK**, then **Next**, followed by **Finish**.
6. After NFS is installed, you can view the details of the **Microsoft Services for NFS** subcomponent to see which of its subcomponents were installed. The subcomponents listed below are preinstalled at the factory. Verify that the Server for NFS and User Name Mapping services have been reinstalled.
   - Microsoft Services for NFS Administration
   - RPC External Data Representation
   - RPC Port Mapper
   - Server for NFS
   - Server for NFS Authentication
   - User Name Mapping

Checklists for cluster server installation

These checklists assist in preparing for installation. Step-by-step instructions begin after the checklists.

Network requirements

- A unique NetBIOS cluster name
- For each node deployed in the cluster the following static IP addresses are required:
  - One for the network adapters on the private network
  - One for the network adapters on the public network
  - One for the virtual server itself

  A single static cluster IP address is required for the entire cluster.

- A domain user account for Cluster service (all nodes must be members of the same domain)
- Each node should have at least two network adapters—one for connection to the public network and the other for the node-to-node private cluster network. If only one network adapter is used for both connections, the configuration is unsupported. A separate private network adapter is required for HCL certification.

Shared disk requirements

NOTE:

Do not allow more than one node access the shared storage devices at the same time until Cluster service is installed on at least one node and that node is online. This can be accomplished through selective storage presentation, SAN zoning, or having only one node online at all times.

- All software components listed in the *HP ProLiant Storage Server SAN Connection and Management* white paper (located on the HP web site at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00663737/c00663737.pdf) must be installed and the fiber cables attached to the HBAs before the cluster installation is started.
- All shared disks, including the Quorum disk, must be accessible from all nodes. When testing connectivity between the nodes and the LUN, only one node should be given access to the LUN at a time.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Cluster installation

During the installation process, nodes are shut down and rebooted. These steps guarantee that the data on disks that are attached to the shared storage bus is not lost or corrupted. This can happen when multiple nodes try to simultaneously write to the same disk that is not yet protected by the cluster software.

Use Table 7 to determine which nodes and storage devices should be presented during each step.

Table 7 Power sequencing for cluster installation

Step	Node 1	Additional Nodes	Storage	Comments
Setting up networks	On	On	Not Presented	Verify that all storage devices on the shared bus are not presented; Power on all nodes.
Setting up shared disks (including the Qurom disk)	On	Off	Presented	Shut down all nodes. Present the shared storage, then power on the first node.
Verifying disk configuration	Off	On	Presented	Shut down first node, power on next node. Repeat this process for all cluster nodes.
Configuring the first node	On	Off	Presented	Shut down all nodes; power on the first node.

Step	Node 1	Additional Nodes	Storage	Comments
Configuring additional nodes	On	On	Presented	Power on the next node after the first node is successfully configured. Complete this process for all cluster nodes.
Post-installation	On	On	Presented	At this point all cluster nodes should be on.

To configure the Cluster service on the storage server, an account must have administrative permissions on each node.

Setting up networks

Verify that all network connections are correct, with private network adapters connected to other private network adapters only, and public network adapters connected to the public network.

Configuring the private network adapter

The following procedures are best practices provided by Microsoft and should be configured on the private network adapter.

- On the **General** tab of the private network adapter, ensure that only TCP/IP is selected.
- Ensure that the **Register this connection's address in DNS** is not selected in the DNS tab under advanced settings for Internet Protocol (TCP/IP) Properties.
- In all cases, set static IP addresses for the private network connector.

Configuring the public network adapter

While the public network adapter's IP address can be automatically obtained if a DHCP server is available, this is not recommended for cluster nodes. HP strongly recommends setting static IP addresses for all network adapters in the cluster, both private and public. If IP addresses are obtained though DHCP, access to cluster nodes could become unavailable if the DHCP server goes down. If DHCP must be used for the public network adapter, use long lease periods to assure that the dynamically assigned lease address remains valid even if the DHCP service is temporarily lost. Keep in mind that Cluster service recognizes only one network interface per subnet.

Renaming the local area connection icons

HP recommends changing the names of the network connections for clarity. The naming helps identify a network and correctly assign its role. For example, “Cluster interconnect” for the private network and “Public connection” for the public network.

Verifying connectivity and name resolution

To verify name resolution, ping each node from a client using the node's machine name instead of its IP address.

Verifying domain membership

All nodes in the cluster must be members of the same domain and able to access a domain controller and a DNS Server.

Setting up a cluster account

The Cluster service requires a domain user account under which the Cluster service can run. This user account must be created before installing Cluster service, because setup requires a user name and password. This user account should be a unique domain account created specifically to administer this cluster. This user account will need to be granted administrator privileges.

About the Quorum disk

HP makes the following Quorum disk recommendations:

- Dedicate a separate disk resource for a Quorum disk. Because the failure of the Quorum disk would cause the entire cluster to fail, HP strongly recommends that the disk resource be a RAID 1 configuration.
- Create a partition with a minimum of 50 megabytes (MB) to be used as a Quorum disk. HP recommends a Quorum disk be 500 MB.

HP recommends assigning the drive letter Q for the Quorum disk. It is also helpful to label the volume Quorum.

NOTE:

It is possible to change the Quorum disk by clicking the Quorum button. This displays a list of available disks that can be used for the Quorum disk. Select the appropriate disk, and then click **OK** to continue.

Configuring shared disks

Use the Windows Disk Management utility to configure additional shared disk resources. Verify that all shared disks are formatted as NTFS and are designated as Basic.

Additional shared disk resources are automatically added into the cluster as physical disk resources during the installation of cluster services.

Verifying disk access and functionality

Write a file to each shared disk resource to verify functionality.

At this time, shut down the first node, power on the next node and repeat the Verifying Disk Access and Functionality step above for all cluster nodes. When it has been verified that all nodes can read and write from the disks, turn off the cluster nodes and power on the first, and then continue with this guide.

Configuring cluster service software

Cluster Administrator (cluadmin) provides the ability to manage, monitor, create and modify clusters and cluster resources.

Using Cluster Administrator

Cluster Administrator shows information about the groups and resources on all of your clusters and specific information about the clusters themselves. A copy of Cluster Administrator is automatically installed on a cluster node when the Cluster service is installed.

Using Cluster Administrator remotely

For remote administration, copies of Cluster Administrator can be installed on other computers on your network. The remote and local copies of Cluster Administrator will be identical. It is also possible to administer an HP ProLiant Storage Server cluster remotely from a computer running Windows NT 4.0 Service Pack 3 or later, Windows 2000 or Windows 2003 using the Cluster Administrator tool.

The HP Storage Server Management Console

Cluster Administrator is available from the HP Storage Server Management Console under the Utilities folder. The HP Storage Server Management Console is accessible using Remote Desktop or a web browser.

Creating a cluster

During the creation of the cluster, Cluster Administrator will analyze and verify the hardware and software configuration and identify potential problems. A comprehensive and easy-to-read report is created, listing any potential configuration issues before the cluster is created.

Some issues that can occur are:

- No shared disk for the Quorum disk. A shared disk must be created with a NTFS partition at least 50 MB in size.
- Use of DHCP addresses for network connections. All Network adapters must be configured with static IP addresses in a cluster configuration.
- File Services for Macintosh and Service for NetWare are not supported in a cluster configuration.
- Dynamic Disks are not supported in a cluster configuration.
- Errors appear on a network adapter that is not configured or does not have an active link. If the network adapter is not going to be used it should be disabled.

Adding nodes to a cluster

Only the Quorum disk should be accessible by the new node while the new node is not a member of the cluster. The new node should not have access to the other LUNs in the cluster until after it has joined the cluster. After the node has joined the cluster, the LUNs may be presented to the new node. Move the physical disk resources over to the new node to confirm functionality.

△ CAUTION:

Presenting other LUNs to the non-clustered system could lead to data corruption.

Geographically dispersed clusters

Cluster nodes can be geographically dispersed to provide an additional layer of fault tolerance. Geographically dispersed clusters are also referred to as stretched clusters.

The following rules must be followed with geographically dispersed clusters:

- A network connection with latency of 500 milliseconds or less ensures that cluster consistency can be maintained. If the network latency is over 500 milliseconds, the cluster consistency cannot be easily maintained.
- All nodes must be on the same subnet.

Cluster groups and resources, including file shares

The Cluster Administrator tool provides complete online help for all cluster administration activities.

Cluster resources include administrative types of resources as well as file shares. The following paragraphs include overview and planning issues for cluster groups, cluster resources, and clustered file shares.

Creating and managing these resources and groups must be managed through Cluster Administrator.

Cluster group overview

A default cluster group is automatically created when the cluster is first created. This default cluster group contains an Internet Protocol (IP) Address resource, a Network Name resource, and the Quorum disk resource. When the new cluster is created, the (IP) address and the cluster name that were specified during setup are set up as the IP address and network name of this default cluster group.

△ CAUTION:

Do not delete or rename the Cluster Group or IP Address. Doing so results in losing the cluster and requires reinstallation of the cluster.

When creating groups, the administrator's first priority is to gain an understanding of how to manage the groups and their resources. Administrators may choose to create a resource group and a virtual server for each node that will contain all resources owned by that node, or the administrator may choose to create a resource group and virtual server for each physical disk resource. Additionally, the administrator should try to balance the load of the groups and their resources on the cluster between the nodes.

Node-based cluster groups

Creating only one resource group and one virtual server for each node facilitates group and resource administration. This setup allows administrators to include all file share resources under one group. Clients access all of the resources owned by one node through a virtual server name.

In node-based cluster groups, each group has its own network name and IP address. The administrator decides on which node to place each physical disk resource. This configuration provides a very coarse level of granularity. All resources within a group must remain on the same node. Only two IP addresses and network names are required. This configuration creates less overhead for resource and network administration. A possible disadvantage of this approach is that the resource groups can potentially grow large when many file shares are created.

Load balancing

The creation of separate cluster groups for each virtual server provides more flexibility in balancing the processing load on the cluster between the two nodes. Each cluster group can be assigned to a cluster node with the preferred owner parameter. For example, if there are two cluster groups, the cluster could be set up to have the first cluster group owned by Node A and the second cluster group owned by Node B. This allows the network load to be handled by both devices simultaneously. If only one cluster group exists, it can only be owned by one node and the other node would not serve any network traffic.

File share resource planning issues

CIFS and NFS are cluster-aware protocols that support the Active/Active cluster model, allowing resources to be distributed and processed on both nodes at the same time. For example, some NFS file share resources can be assigned to a group owned by a virtual server for Node A and additional NFS file share resources can be assigned to a group owned by a virtual server for Node B.

Configuring the file shares as cluster resources provides for high availability of file shares. Because the resources are placed into groups, ownership of the files can easily move from one node to the other, as circumstances require. If the cluster node owning the group of file shares should be shut down or fail, the other node in the cluster will begin sharing the directories until the original owner node is brought back on line. At that time, ownership of the group and its resources can be brought back to the original owner node.

Resource planning

1. Create a cluster group for each node in the cluster with an IP address resource and a network name resource.

   Cluster resource groups are used to balance the processing load on the servers. Distribute ownership of the groups between the virtual servers.

2. For NFS environments, configure the NFS server.

   NFS specific procedures include entering audit and file lock information as well as setting up client groups and user name mappings. These procedures are not unique to a clustered deployment and are detailed in the Microsoft Services for NFS section within the “Other network file and print services” chapter. Changes to NFS setup information are automatically replicated to all nodes in a cluster.

3. Create the file share resources.

4. Assign ownership of the file share resources to the resource groups.

   a. Divide ownership of the file share resource between the resource groups, which are in turn distributed between the virtual servers, for effective load balancing.
   b. Verify that the physical disk resource for this file share is also included in this group.
   c. Verify that the resources are dependent on the virtual servers and physical disk resources from which the file share was created.

Permissions and access rights on share resources

File Share and NFS Share permissions must be managed using the Cluster Administrator tool versus the individual shares on the file system themselves via Windows Explorer. Administering them through the Cluster Administrator tool allows the permissions to migrate from one node to other. In addition, permissions established using Explorer are lost after the share is failed or taken offline.

NFS cluster-specific issues

For convenience, all suggestions are listed below:

- Back up user and group mappings.

  To avoid loss of complex advanced mappings in the case of a system failure, back up the mappings whenever the mappings have been edited or new mappings have been added.

- Map consistently.

  Groups that are mapped to each other should contain the same users and the members of the groups should be properly mapped to each other to ensure proper file access.

- Map properly.
  - Valid UNIX users should be mapped to valid Windows users.
  - Valid UNIX groups should be mapped to valid Windows groups.
  - Mapped Windows user must have the “Access this computer from the Network privilege” or the mapping will be squashed.
  - The mapped Windows user must have an active password, or the mapping will be squashed.

- In a clustered deployment, create user name mappings using domain user accounts.

  Because the security identifiers of local accounts are recognized only by the local server, other nodes in the cluster will not be able to resolve those accounts during a failover. Do not create mappings using local user and group accounts.

- In a clustered deployment, administer user name mapping on a computer that belongs to a trusted domain.

  If NFS administration tasks are performed on a computer that belongs to a domain that is not trusted by the domain of the cluster, the changes are not properly replicated among the nodes in the cluster.

- In a clustered deployment, if PCNFS password and group files are being used to provide user and group information, these files must be located on each node of the system.

  Example: If the password and group files are located at `c:\maps` on node 1, then they must also be at `c:\maps` on node 2. The contents of the password and group files must be the same on both nodes as well.

  These password and group files on each server node must be updated periodically to maintain consistency and prevent users or groups from being inadvertently squashed.

Non cluster aware file sharing protocols

Services for Macintosh (SFM), File and Print Services for NetWare, HTTP file sharing protocols are not cluster aware and will experience service interruption if installed on a clustered resource during failover events of the resource. Service interruptions will be similar to those experienced during a server outage. Data that has not been saved to disk prior to the outage will experience data loss. In the case of SFM, it is not supported because SFM maintains state information in memory. Specifically, the Macintosh volume index is located in paged pool memory. Using SFM in clustered mode is not supported and may result in data loss similar in nature to a downed server should the resource it is based on fails over to the opposing node.

Adding new storage to a cluster

Present the new storage to one node in the cluster. This can be accomplished through selective storage presentation or through SAN zoning.

The tasks described below are used to add storage to a cluster. See the online help for clustering for additional details.

Creating physical disk resources

A physical disk resource must reside within a cluster group. An existing cluster group can be used or a new cluster group must be created. For information on creating disk resources, see the cluster online help topic *Physical Disk resource type*.

NOTE:

- Physical disk resources usually do not have any dependencies set.
- In multi-node clusters it is necessary to specify the node to move the group to. When a cluster group is moved to another node, all resources in that group are moved.
- When a physical disk resource is owned by a node, the disk appears as an unknown, unreadable disk to all other cluster nodes. This is a normal condition. When the physical disk resource moves to another node, the disk resource then becomes readable.

Creating file share resources

To create a file share resource, see two clustering online help topics:

- Create a cluster-managed file share
- Using a server cluster with large numbers of file shares

NOTE:

- A file share resource must reside in the same cluster group as the physical disk resource it will reside on.
- The physical disk resource specified in this step must reside in the same cluster group as specified in the beginning of this wizard.

Creating NFS share resources

To create an NFS share resource, see "MSNFS administration on a server cluster" on page 106.

Shadow copies in a cluster

It is recommended that the location of the cache file be placed on a separate disk from the original data. In this case, a physical disk resource for the cache file disk should be created in the same cluster group as the intended Shadow Copy resource and the volume for which snapshots will be enabled. The resource should be created prior to the establishment of Shadow Copies. The Shadow Copy resource should be dependent on both the original physical disk resource and the physical disk resource that contains the cache file.

For more information, see the following topics in the clustering online help:

- Using Shadow Copies of Shared Folders in a server cluster
- Enable Shadow Copies for shared folders in a cluster

Extend a LUN in a cluster

To extend a LUN on a storage array in a cluster, review the requirements and procedures from the storage array hardware provider for expanding or extending storage.

For additional information associated with extending a LUN in a cluster, see the following Microsoft Knowledge Base articles:

- How to extend the partition of a cluster shared disk
  http://support.microsoft.com/default.aspx?scid=kb;en-us;304736
- How to replace a disk that is in a cluster and use of the Cluster Recovery utility
  http://support.microsoft.com/kb/305793

MSNFS administration on a server cluster

The Microsoft Services for Network File System (NFS) online help provides server cluster information for the following topics:

- Configuring shared folders on a server cluster
  - Configuring an NFS share as a cluster resource
  - Modifying an NFS shared cluster resource
  - Deleting an NFS shared cluster resource
- Using Microsoft Services for NFS with server clusters
  - Understanding how Server for NFS works with server clusters
  - Using Server for NFS on a server cluster
- Configuring User Name Mapping on a server cluster

For further details, see the online help for Microsoft Services for Network File System.

Best practices for running Server for NFS in a server cluster

- Stop Server for NFS before stopping the server cluster.
- Ensure share availability when a node fails.
- Use the appropriate tool to manage Network File System (NFS) share cluster resources.
- Avoid conflicting share names.
- Ensure the availability of audit logs.
- Move file shares or take them offline before stopping Server for NFS.
- Take resources offline before modifying.
- Administer Server for NFS only from computers in a trusted domain.
- Restart the Server for NFS service after the cluster service restarts.
- Choose the appropriate sharing mode.
- Use the command line properly when creating or modifying NFS share cluster resources.
- Use hard mounts.
- Use the correct virtual server name.

Print services in a cluster

The Windows Server 2003 Cluster service implementation increases availability of critical print servers. A print spooler service on a clustered print server may be hosted on any of the nodes in the cluster. As with all cluster resources, clients should access the print server by its virtual network name or virtual IP address.

Creating a cluster printer spooler

Printer spoolers should be created in a separate group dedicated to this purpose for ease of management. For each printer spooler, a physical resource is required to instantiate the print spooler resource. In some cases, dedicated physical resources are not available and hence sharing of the physical resource among other members of the group is acceptable, remembering that all members of a group are managed as a unit. Hence, the group will failover and failback as a group.

To create a printer spooler:

1. Create a dedicated group (if desired).
2. Create a physical resource (disk) (if required, see note).
3. Create an IP address resource for the Virtual Server to be created (if required, see note).
4. Create a Virtual Server Resource (Network Name) (if required, see note).

NOTE:

If the printer spool resource is added to an existing group with a physical resource, IP address, and virtual server resource, steps 1-4 are not required.

5. Create a Print Spool resource.
6. To add a printer to the virtual server:
   a. Double-click the printers and faxes icon.
   b. Right-click the new screen, and then click **add printer**. A wizard starts.
   c. Click **create a new port**, and then click **Next**.
   d. Enter the IP address of the network printer.
   e. Update the Port Name if desired, click **Next**, and then click **Finish**.
   f. Select the appropriate driver, and then click **Next**.
   g. If presented with a dialog to replace the driver present, click **keep the driver**, and then click **Next**.
   h. Name the printer, and then click **Next**.
   i. Provide a share name for the printer for network access, and then click **Next**.
   j. Provide location information and comments, and then click **Next**.
   k. Click **Yes** to print a test page, click **Next**, and then click **Finish**.
   l. A dialog box appears regarding the test page. Select the appropriate answer.

The Printer Spool is now a clustered resource.

Advanced cluster administration procedures

Failing over and failing back

As previously mentioned, when a node goes offline, all resources dependent on that node are automatically failed over to another node. Processing continues, but in a reduced manner, because all operations must be processed on the remaining node(s). In clusters containing more than two nodes, additional fail over rules can be applied. For instance, groups can be configured to fail over different nodes to balance the additional work load imposed by the failed node. Nodes can be excluded from the possible owners list to prevent a resource from coming online on a particular node. Lastly the preferred owners list can be ordered, to provide an ordered list of failover nodes. Using these tools, the failover of resources can be controlled with in a multinode cluster to provide a controlled balanced failover methodology that balances the increased work load.

Because operating environments differ, the administrator must indicate whether the system will automatically fail the resources (organized by resource groups) back to their original node or will leave the resources failed over, waiting for the resources to be moved back manually.

NOTE:

If the storage server is not set to automatically fail back the resources to their designated owner, the resources must be moved back manually each time a failover occurs.

Restarting one cluster node

CAUTION:

Restarting a cluster node should be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being restarted. Attached connections can be viewed through the Management Console on the storage server Desktop using Terminal Services. From the Management Console, select **File Sharing > Shared Folders > Sessions**.

The physical process of restarting one of the nodes of a cluster is the same as restarting a storage server in single node environment. However, additional caution is needed.

Restarting a cluster node causes all cluster resources served by that node to fail over to the other nodes in the cluster based on the failover policy in place. Until the failover process completes, any currently executing read and write operations will fail. Other node(s) in the cluster will be placed under a heavier load by the extra work until the restarted node comes up and the resources are moved back.

Shutting down one cluster node

CAUTION:

Shutting down a cluster node must be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being shutdown.

Shutting down a cluster node causes all cluster resources served by that node to fail over to the other node(s). This causes any currently executing client read and write operations to fail until the cluster failover process completes. The other node(s) are placed under a heavier load by the extra work until the second node is powered up and rejoins the cluster.

Powering down the cluster

The power down process for the storage server cluster is similar to the process for a single node, but with the cluster, extra care must be taken with the storage subsystem and the sequence of the shutdown.

The power down process is divided into two main steps:

1. Shutting down the cluster nodes
2. Removing power from the cluster nodes

The sequence of these steps is critical. The devices must be shut down before the storage subsystem. Improperly shutting down the nodes and the storage subsystem causes corruption and loss of data.

△ CAUTION:

Before powering down the cluster nodes, follow the proper shutdown procedure as previously illustrated. See "Shutting down one cluster node." Only one cluster node should be shut down at a time.

Powering up the cluster

The power up process for the storage server cluster is more complex than it is for a single node because extra care must be taken with the storage subsystem.

The sequence of the power up steps is critical. Improper power up procedures can cause corruption and loss of data.

△ CAUTION:

Do not power up the cluster nodes without first powering up the storage subsystem, and verifying it is operating normally.

Nodes should be powered up separately allowing one node to form the cluster prior to powering up the additional node(s). To power up the cluster nodes:

1. After the storage subsystem is confirmed to be operating normally, power up a single node. Wait for the node to come completely up before powering up the subsequent node(s).

   If more than one node is powered up at the same time, the first node that completes the sequence gains ownership of the cluster quorum and controls the cluster database. Designate a particular node as the usual cluster quorum owner by always powering up that node first and letting it completely restart before powering up additional cluster node(s).
2. Power up the additional cluster node(s). Each node should be allowed to start fully, prior to starting a subsequent node.

Additional information and references for cluster services

The following web sites provide detailed information for clustered environments for Windows Server 2003, which also applies to Windows Storager Server 2003:

- Cluster services
  http://www.microsoft.com/windowsserver2003/technologies/clustering/default.mspx
- How to: Set up a clustered print server
  http://support.microsoft.com/default.aspx?scid=kb;en-us;278455
- How to: Set up a print spooler on Microsoft Cluster Server
  http://support.microsoft.com/kb/197046/
- How to: Troubleshoot printing issues on a Windows Server 2003 Cluster
  http://support.microsoft.com/default.aspx?scid=kb;en-us;302539
- Creating and configuring a highly available print server under Microsoft Windows Server 2003 using a server cluster
  http://www.microsoft.com/WindowsServer2003/techinfo/overview/availableprinter.mspx

8 Troubleshooting, servicing, and maintenance

Troubleshooting the storage server

The "Support and troubleshooting" task at the HP Support & Drivers web site (http://www.hp.com/go/support) can be used to troubleshoot problems with the storage server. After entering the storage server name and designation (for example, ML110 G5 storage server) or component information (for example, Array Configuration Utility), use the following links for troubleshooting information:

- Download drivers and software—This area provides drivers and software for your operating system.
- Troubleshoot a problem—This area provides a listing of customer notices, advisories, and bulletins applicable for the product or component.
- Manuals—This area provides the latest user documentation applicable to the product or component. User guides can be a useful source for troubleshooting information. For most storage server hardware platforms, the following ProLiant server manuals may be useful for troubleshooting assistance:
  - **HP ProLiant <model> Server User Guide or HP ProLiant <model> Server Maintenance and Service Guide** (where <model> is the product model of the storage server, such as ML110 G5).

    These guides contain specific troubleshooting information for the server. The guides are available by selecting the applicable ProLiant Server series model, then the Manuals (guides, supplements, addendums, etc.) link.

    For example, instead of using "ML110 G5 storage server," enter "ML110 G5 server" for the product to search, then select the "HP ProLiant ML110 Server series" link, then the Manuals (guides, supplements, addendums, etc.) link to locate the guide.
  - **HP ProLiant Servers Troubleshooting Guide**

    The guide provides common procedures and solutions for many levels of troubleshooting with a ProLiant server. The guide is available at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00300504/c00300504.pdf.

IMPORTANT:

Not all troubleshooting procedures found in ProLiant server guides may apply to the ProLiant Storage Server. If necessary, check with your HP Support representative for further assistance.

For software related components and issues, online help or user guide documentation may offer troubleshooting assistance. The release notes for the storage server product line is updated frequently. The document contains issues and workarounds to a number of categories for the storage servers.

Known issues and workarounds for the storage server products and the service release are addressed in release notes. To view the latest release notes, go to http://www.hp.com/support/manuals. Under the storage section, click **NAS** and then select your product.

WEBES (Web Based Enterprise Services)

WEBES is a tool suite aimed at preventing or reducing your system's down time. The tool suite has the following components:

- CCAT (Computer Crash Analysis Tool)
- SEA (System Event Analyzer)

If you have a warranty or service contract with HP you are entitled to these tools free of charge. You must, however, upgrade the tools at least once a year because the software expires after one year. For more information about WEBES, see http://h18023.www1.hp.com/support/svctools/webes/.

To install WEBES on your storage server, run the setup executable located in the `C:\hpnas\Components\WEBES` folder.

Maintenance and service

HP provides specific documentation for maintaining and servicing your storage server and offers a customer self repair program.

Maintenance and service documentation

For specific documentation for the maintenance and servicing of HP ProLiant Storage Servers, see the *HP ProLiant <model> Server Maintenance and Service Guide* for your storage server model. This document can be obtained at http://www.hp.com/support/manuals. Under the servers section, select **ProLiant and tc series servers**, and then select your product.

Additional documentation can also be found on the inside of the access panel of certain server models.

Maintenance updates

Regular updates to the storage server are supplied on the HP ProLiant Storage Server Service Release DVD. The Service Release DVD can be obtained at http://www.software.hp.com.

Individual updates for each product are available for download from the HP Support web site at http://h18023.www1.hp.com/support/selfrepair/na/replace_part.asp.

System updates

System updates to the hardware (BIOS, firmware, drivers), critical updates, and hotfixes for the operating system and other related software updates are bundled on the Service Release DVD.

Firmware updates

Firmware is software that is stored in Read-Only Memory (ROM). Firmware is responsible for the behavior of the system when it is first switched on and for passing control of the server to the operating system. When referring to the firmware on the system board of the server, it is called the System ROM or the BIOS. When referring to the firmware on another piece of hardware configured in the server, it is called Option ROM. ProLiant servers have hard drives, Smart Array Controllers, Remote Insight

Lights-Out Edition (RILOE), Remote Insight Lights-Out Edition II (RILOE II) and Integrated Lights-Out options that have firmware that can be updated.

It is important to update the firmware (also called "flashing the ROM") as part of regular server maintenance. In addition, checking for specific firmware updates in between regular updates helps to keep the server performing optimally. HP recommends checking for a firmware update before sending a part back to HP for replacement.

Certificate of Authenticity

The Certificate of Authenticity (COA) label is used to:

- Upgrade the factory-installed operating system using the Microsoft Upgrade program for license validation.
- Reinstall the operating system because of a failure that has permanently disabled it.

The COA label location varies by server model. On rack-mounted server models, the COA label is located either on the front section of the right panel or on the right front corner of the top panel. On tower models, the COA label is located toward the rear of the top panel of the server.

9 System recovery

This chapter describes how to use the System Recovery DVD that is provided with your storage server.

The System Recovery DVD

The HP ProLiant Storage Server System Recovery DVD that is provided with your storage server allows you to install an image or recover from a catastrophic failure.

At any later time, you may boot from the DVD and restore the server to the factory condition. This allows you to recover the system if all other means to boot the server fail.

While the recovery process makes every attempt to preserve the existing data volumes, you should have a backup of your data if at all possible before recovering the system.

To restore a factory image

1. Insert the System Recovery DVD. The main window appears.
2. Choose **Restore Factory Image**.

Systems with a DON'T ERASE partition

The DON'T ERASE logical disk supports the restoration process only and does not host a secondary operating system. Be sure to back up your user data, and then use the Recovery and Installation DVD to restore the server to the factory state.

Managing disks after a restoration

After a system has been restored, drive letters may be assigned to the wrong volume. Windows Storage Server 2003 assigns drive letters after the restoration in the order of discovery. To help maintain drive letter information, placing the drive letter into a volume label is recommended. To change the drive letters to the appropriate one, go into Disk Management and perform the following steps for each volume:

1. Right-click the volume that needs to be changed.
2. Select **Change drive Letter and Paths**.
3. In the **Change drive Letter and Paths** dialog box, select **Change**.
4. Select the appropriate drive letter, then click **OK**.
5. Click **Yes** to confirm the drive letter change.
6. Click **Yes** to continue. If the old drive letter needs to be re-used, reboot the server after clicking Yes.

A Regulatory compliance and safety

Federal Communications Commission notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that may reasonably be expected to be installed in a business or commercial environment. Class B devices are those that may reasonably be expected to be installed in a residential environment (personal computers, for example). The FCC requires devices in both classes to bear a label indicating the interference potential of the device as well as additional operating instructions for the user.

The rating label on the device shows which class (A or B) the equipment falls into. Class B devices have an FCC logo or FCC ID on the label. Class A devices do not have an FCC logo or FCC ID on the label. Once the class of the device is determined, see the following corresponding statement.

Class A equipment

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Class B equipment

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or television technician for help.

Declaration of conformity for products marked with the FCC logo, United States only

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

For questions regarding your product, contact:

Hewlett-Packard Company

P. O. Box 692000, Mail Stop 530113

Houston, Texas 77269-2000

Or, call

1-800- 652-6672

For questions regarding this FCC declaration, contact:

Hewlett-Packard Company

P. O. Box 692000, Mail Stop 510101

Houston, Texas 77269-2000

Or, call

(281) 514-3333

To identify this product, see the Part, Series, or Model number found on the product.

Modifications

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company may void the user's authority to operate the equipment.

Cables

Connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to maintain compliance with FCC Rules and Regulations.

Laser compliance

This product may be provided with an optical storage device (that is, CD or DVD drive) and/or fiber optic transceiver. Each of these devices contains a laser that is classified as a Class 1 Laser Product in accordance with US FDA regulations and the IEC 60825–1. The product does not emit hazardous laser radiation.

⚠ WARNING!

Use of controls or adjustments or performance of procedures other than those specified herein or in the installation guide of the laser product may result in hazardous radiation exposure. To reduce the risk of exposure to hazardous radiation:

- Do not try to open the module enclosure. There are no user-serviceable components inside.
- Do not operate controls, make adjustments, or perform procedures to the laser device other than those specified herein.
- Allow only HP authorized service technicians to repair the unit.

The Center for Devices and Radiological Health (CDRH) of the U.S. Food and Drug Administration implemented regulations for laser products on August 2, 1976. These regulations apply to laser products manufactured from August 1, 1976. Compliance is mandatory for products marketed in the United States.

International notices and statements

Canadian notice (Avis Canadien)

Class A equipment

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Class B equipment

This Class B digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European Union notice

CE Products bearing the CE marking comply with the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community and if this product has telecommunication functionality, the R&TTE Directive (1999/5/EC).

Compliance with these directives implies conformity to the following European Norms (in parentheses are the equivalent international standards and regulations):

- EN 55022 (CISPR 22) - Electromagnetic Interference
- EN55024 (IEC61000-4-2, 3, 4, 5, 6, 8, 11) - Electromagnetic Immunity
- EN61000-3-2 (IEC61000-3-2) - Power Line Harmonics
- EN61000-3-3 (IEC61000-3-3) - Power Line Flicker
- EN 60950 (IEC 60950) - Product Safety

BSMI notice

警告使用者:
這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Japanese notice

ご使用になっている装置にVCCIマークが付いていましたら、次の説明文をお読み下さい。

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスB情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。
取扱説明書に従って正しい取り扱いをして下さい。

VCCIマークが付いていない場合には、次の点にご注意下さい。

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です　この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean notice A&B

Class A equipment

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Class B equipment

B급 기기 (가정용 정보통신기기)

이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거지역에서는 물론 모든지역에서 사용할 수 있습니다.

Safety

Battery replacement notice

⚠ WARNING!

The computer contains an internal lithium manganese dioxide, a vanadium pentoxide, or an alkaline battery pack. A risk of fire and burns exists if the battery pack is not properly handled. To reduce the risk of personal injury:

- Do not attempt to recharge the battery.
- Do not expose the battery to temperatures higher than 60°C (140°F).
- Do not disassemble, crush, puncture, short external contacts, or dispose of in fire or water.

 Batteries, battery packs, and accumulators should not be disposed of together with the general household waste. To forward them to recycling or proper disposal, please use the public collection system or return them to HP, an authorized HP Partner, or their agents.

For more information about battery replacement or proper disposal, contact an authorized reseller or an authorized service provider.

Taiwan battery recycling notice

The Taiwan EPA requires dry battery manufacturing or importing firms in accordance with Article 15 of the Waste Disposal Act to indicate the recovery marks on the batteries used in sales, giveaway or promotion. Contact a qualified Taiwanese recycler for proper battery disposal.

Power cords

The power cord set must meet the requirements for use in the country where the product was purchased. If the product is to be used in another country, purchase a power cord that is approved for use in that country.

The power cord must be rated for the product and for the voltage and current marked on the product electrical rating label. The voltage and current rating of the cord should be greater than the voltage and current rating marked on the product. In addition, the diameter of the wire must be a minimum of 1.00 mm^2 or 18 AWG, and the length of the cord must be between 1.8 m (6 ft) and 3.6 m (12 ft). If you have questions about the type of power cord to use, contact an HP authorized service provider.

NOTE:

Route power cords so that they will not be walked on and cannot be pinched by items placed upon or against them. Pay particular attention to the plug, electrical outlet, and the point where the cords exit from the product.

Japanese power cord notice

製品には、同梱された電源コードをお使い下さい。
同梱された電源コードは、他の製品では使用出来ません。

Electrostatic discharge

To prevent damage to the system, be aware of the precautions you need to follow when setting up the system or handling parts. A discharge of static electricity from a finger or other conductor may damage system boards or other static-sensitive devices. This type of damage may reduce the life expectancy of the device.

Preventing electrostatic discharge

To prevent electrostatic damage, observe the following precautions:

- Avoid hand contact by transporting and storing products in static-safe containers.
- Keep electrostatic-sensitive parts in their containers until they arrive at static-free workstations.
- Place parts on a grounded surface before removing them from their containers.
- Avoid touching pins, leads, or circuitry.
- Always be properly grounded when touching a static-sensitive component or assembly.

Grounding methods

There are several methods for grounding. Use one or more of the following methods when handling or installing electrostatic-sensitive parts:

- Use a wrist strap connected by a ground cord to a grounded workstation or computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm ± 10 percent resistance in the ground cords. To provide proper grounding, wear the strap snug against the skin.
- Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.
- Use conductive field service tools.
- Use a portable field service kit with a folding static-dissipating work mat.

If you do not have any of the suggested equipment for proper grounding, have an authorized reseller install the part.

NOTE:

For more information on static electricity, or for assistance with product installation, contact your authorized reseller.

Waste Electrical and Electronic Equipment (WEEE) directive

Czechoslovakian notice

Likvidace za ízení soukromými domácími uživateli v Evropské unii

Tento symbol na produktu nebo balení ozna uje výrobek, který nesmí být vyhozen spolu s ostatním domácím odpadem. Povinností uživatele je p edat takto ozna ený odpad na p edem ur ené sb rné místo pro recyklaci elektrických a elektronických za ízení. Okamžité t íd ní a recyklace odpadu pom že uchovat p írodní prost edí a zajistí takový zp sob recyklace, který ochrání zdraví a životní prost edí lov ka. Další informace o možnostech odevzdání odpadu k recyklaci získáte na p íslušném obecním nebo m stském ú ad , od firmy zabývající se sb rem a svozem odpadu nebo v obchod , kde jste produkt zakoupili.

Danish notice

Bortskaffelse af affald fra husstande i den Europæiske Union

Hvis produktet eller dets emballage er forsynet med dette symbol, angiver det, at produktet ikke må bortskaffes med andet almindeligt husholdningsaffald. I stedet er det dit ansvar at bortskaffe kasseret udstyr ved at aflevere det på den kommunale genbrugsstation, der forestår genvinding af kasseret elektrisk og elektronisk udstyr. Den centrale modtagelse og genvinding af kasseret udstyr i forbindelse med bortskaffelsen bidrager til bevarelse af naturlige ressourcer og sikrer, at udstyret genvindes på en måde, der beskytter både mennesker og miljø. Yderligere oplysninger om, hvor du kan aflevere kasseret udstyr til genvinding, kan du få hos kommunen, den lokale genbrugsstation eller i den butik, hvor du købte produktet.

Dutch notice

Verwijdering van afgedankte apparatuur door privé-gebruikers in de Europese Unie

Dit symbool op het product of de verpakking geeft aan dat dit product niet mag worden gedeponeerd bij het normale huishoudelijke afval. U bent zelf verantwoordelijk voor het inleveren van uw afgedankte apparatuur bij een inzamelingspunt voor het recyclen van oude elektrische en elektronische apparatuur. Door uw oude apparatuur apart aan te bieden en te recyclen, kunnen natuurlijke bronnen worden behouden en kan het materiaal worden hergebruikt op een manier waarmee de volksgezondheid en het milieu worden beschermd. Neem contact op met uw gemeente, het afvalinzamelingsbedrijf of de winkel waar u het product hebt gekocht voor meer informatie over inzamelingspunten waar u oude apparatuur kunt aanbieden voor recycling.

English notice

Disposal of waste equipment by users in private household in the European Union

This symbol on the product or on its packaging indicates that this product must not be disposed of with your other household waste. Instead, it is your responsibility to dispose of your waste equipment by handing it over to a designated collection point for recycling of waste electrical and electronic equipment. The separate collection and recycling of your waste equipment at the time of disposal will help to conserve natural resources and ensure that it is recycled in a manner that protects human health and the environment. For more information about where you can drop off your waste equipment for recycling, please contact your local city office, your household waste disposal service, or the shop where you purchased the product.

Estonian notice

Seadmete jäätmete kõrvaldamine eramajapidamistes Euroopa Liidus

See tootel või selle pakendil olev sümbol näitab, et kõnealust toodet ei tohi koos teiste majapidamisjäätmetega kõrvaldada. Teie kohus on oma seadmete jäätmed kõrvaldada, viies need elektri- ja elektroonikaseadmete jäätmete ringlussevõtmiseks selleks ettenähtud kogumispunkti. Seadmete jäätmete eraldi kogumine ja ringlussevõtmine kõrvaldamise ajal aitab kaitsta loodusvarasid ning tagada, et ringlussevõtmine toimub viisil, mis kaitseb inimeste tervist ning keskkonda. Lisateabe saamiseks selle kohta, kuhu oma seadmete jäätmed ringlussevõtmiseks viia, võtke palun ühendust oma kohaliku linnakantselei, majapidamisjäätmete kõrvaldamise teenistuse või kauplusega, kust Te toote ostsite.

Finnish notice

Laitteiden hävittäminen kotitalouksissa Euroopan unionin alueella

Jos tuotteessa tai sen pakkauksessa on tämä merkki, tuotetta ei saa hävittää kotitalousjätteiden mukana. Tällöin hävitettävä laite on toimitettava sähkölaitteiden ja elektronisten laitteiden kierrätyspisteeseen. Hävitettävien laitteiden erillinen käsittely ja kierrätys auttavat säästämään luonnonvaroja ja varmistamaan, että laite kierrätetään tavalla, joka estää terveyshaitat ja suojelee luontoa. Lisätietoja paikoista, joihin hävitettävät laitteet voi toimittaa kierrätettäväksi, saa ottamalla yhteyttä jätehuoltoon tai liikkeeseen, josta tuote on ostettu.

French notice

Élimination des appareils mis au rebut par les ménages dans l'Union européenne

Le symbole apposé sur ce produit ou sur son emballage indique que ce produit ne doit pas être jeté avec les déchets ménagers ordinaires. Il est de votre responsabilité de mettre au rebut vos appareils en les déposant dans les centres de collecte publique désignés pour le recyclage des équipements électriques et électroniques. La collecte et le recyclage de vos appareils mis au rebut indépendamment du reste des déchets contribue à la préservation des ressources naturelles et garantit que ces appareils seront recyclés dans le respect de la santé humaine et de l'environnement. Pour obtenir plus d'informations sur les centres de collecte et de recyclage des appareils mis au rebut, veuillez contacter les autorités locales de votre région, les services de collecte des ordures ménagères ou le magasin dans lequel vous avez acheté ce produit.

German notice

Entsorgung von Altgeräten aus privaten Haushalten in der EU

Das Symbol auf dem Produkt oder seiner Verpackung weist darauf hin, dass das Produkt nicht über den normalen Hausmüll entsorgt werden darf. Benutzer sind verpflichtet, die Altgeräte an einer Rücknahmestelle für Elektro- und Elektronik-Altgeräte abzugeben. Die getrennte Sammlung und ordnungsgemäße Entsorgung Ihrer Altgeräte trägt zur Erhaltung der natürlichen Ressourcen bei und

garantiert eine Wiederverwertung, die die Gesundheit des Menschen und die Umwelt schützt. Informationen dazu, wo Sie Rücknahmestellen für Ihre Altgeräte finden, erhalten Sie bei Ihrer Stadtverwaltung, den örtlichen Müllentsorgungsbetrieben oder im Geschäft, in dem Sie das Gerät erworben haben.

Greek notice

Απόρριψη άχρηστου εξοπλισμού από χρήστες σε ιδιωτικά νοικοκυριά στην Ευρωπαϊκή Ένωση

Το σύμβολο αυτό στο προϊόν ή τη συσκευασία του υποδεικνύει ότι το συγκεκριμένο προϊόν δεν πρέπει να διατίθεται μαζί με τα άλλα οικιακά σας απορρίμματα. Αντίθετα, είναι δική σας ευθύνη να απορρίψετε τον άχρηστο εξοπλισμό σας παραδίδοντάς τον σε καθορισμένο σημείο συλλογής για την ανακύκλωση άχρηστου ηλεκτρικού και ηλεκτρονικού εξοπλισμού. Η ξεχωριστή συλλογή και ανακύκλωση του άχρηστου εξοπλισμού σας κατά την απόρριψη θα συμβάλει στη διατήρηση των φυσικών πόρων και θα διασφαλίσει ότι η ανακύκλωση γίνεται με τρόπο που προστατεύει την ανθρώπινη υγεία και το περιβάλλον. Για περισσότερες πληροφορίες σχετικά με το πού μπορείτε να παραδώσετε τον άχρηστο εξοπλισμό σας για ανακύκλωση, επικοινωνήστε με το αρμόδιο τοπικό γραφείο, την τοπική υπηρεσία διάθεσης οικιακών απορριμμάτων ή το κατάστημα όπου αγοράσατε το προϊόν.

Hungarian notice

Készülékek magánháztartásban történ selejtezése az Európai Unió területén

A készüléken, illetve a készülék csomagolásán látható azonos szimbólum annak jelzésére szolgál, hogy a készülék a selejtezés során az egyéb háztartási hulladéktól eltér módon kezelend . A vásárló a hulladékká vált készüléket köteles a kijelölt gy jt helyre szállítani az elektromos és elektronikai készülékek újrahasznosítása céljából. A hulladékká vált készülékek selejtezéskori begy jtése és újrahasznosítása hozzájárul a természeti er források meg rzéséhez, valamint biztosítja a selejtezett termékek környezetre és emberi egészségre nézve biztonságos feldolgozását. A begy jtés pontos helyér l b vebb tájékoztatást a lakhelye szerint illetékes önkormányzattól, az illetékes szemételtakarító vállalattól, illetve a terméket elárusító helyen kaphat.

Italian notice

Smaltimento delle apparecchiature da parte di privati nel territorio dell'Unione Europea

Questo simbolo presente sul prodotto o sulla sua confezione indica che il prodotto non può essere smaltito insieme ai rifiuti domestici. È responsabilità dell'utente smaltire le apparecchiature consegnandole presso un punto di raccolta designato al riciclo e allo smaltimento di apparecchiature elettriche ed elettroniche. La raccolta differenziata e il corretto riciclo delle apparecchiature da smaltire permette di proteggere la salute degli individui e l'ecosistema. Per ulteriori informazioni relative ai punti di raccolta delle apparecchiature, contattare l'ente locale per lo smaltimento dei rifiuti, oppure il negozio presso il quale è stato acquistato il prodotto.

Latvian notice

Nolietotu iek rtu izn cin šanas noteikumi lietot jiem Eiropas Savien bas priv taj s m jsaimniec b s

Š ds simbols uz izstr d juma vai uz t iesai ojuma nor da, ka šo izstr d jumu nedr kst izmest kop ar citiem sadz ves atkritumiem. J s atbildat par to, lai nolietot s iek rtas tiktu nodotas speci li iek rtotos punktos, kas paredz ti izmantoto elektrisko un elektronisko iek rtu sav kšanai otrreiz jai p rstr dei. Atseviš a nolietoto iek rtu sav kšana un otrreiz j p rstr de pal dz s saglab t dabas resursus un garant s, ka š s iek rtas tiks otrreiz ji p rstr d tas t d veid , lai pasarg tu vidi un cilv ku vesel bu. Lai uzzin tu, kur nolietot s iek rtas var izmest otrreiz jai p rstr dei, j v ršas savas dz ves vietas pašvald b , sadz ves atkritumu sav kšanas dienest vai veikal , kur izstr d jums tika nopirkts.

Lithuanian notice

Vartotoj iš priva i nam ki rangos atliek šalinimas Europos S jungoje

Šis simbolis ant gaminio arba jo pakuot s rodo, kad šio gaminio šalinti kartu su kitomis nam kio atliekomis negalima. Šalintinas rangos atliekas privalote pristatyti speciali surinkimo viet elektros ir elektronin s rangos atliekoms perdirbti. Atskirai surenkamos ir perdirbamos šalintinos rangos atliekos pad s saugoti gamtinius išteklius ir užtikrinti, kad jos bus perdirbtos tokiu b du, kuris nekenkia žmoni sveikatai ir aplinkai. Jeigu norite sužinoti daugiau apie tai, kur galima pristatyti perdirbtinas rangos atliekas, kreipkit s savo seni nij , nam kio atliek šalinimo tarnyb arba parduotuv , kurioje sigijote gamin .

Polish notice

Pozbywanie si zu ytego sprz tu przez u ytkowników w prywatnych gospodarstwach domowych w Unii Europejskiej

Ten symbol na produkcie lub jego opakowaniu oznacza, e produktu nie wolno wyrzuca do zwykłych pojemników na mieci. Obowi zkiem u ytkownika jest przekazanie zu ytego sprz tu do wyznaczonego punktu zbiórki w celu recyklingu odpadów powstałych ze sprz tu elektrycznego i elektronicznego. Osobna zbiórka oraz recykling zu ytego sprz tu pomog w ochronie zasobów naturalnych i zapewni ponowne wprowadzenie go do obiegu w sposób chroni cy zdrowie człowieka i rodowisko. Aby uzyska wi cej informacji o tym, gdzie mo na przekaza zu yty sprz t do recyklingu, nale y si skontaktowa z urz dem miasta, zakładem gospodarki odpadami lub sklepem, w którym zakupiono produkt.

Portuguese notice

Descarte de Lixo Elétrico na Comunidade Européia

Este símbolo encontrado no produto ou na embalagem indica que o produto não deve ser descartado no lixo doméstico comum. É responsabilidade do cliente descartar o material usado (lixo elétrico), encaminhando-o para um ponto de coleta para reciclagem. A coleta e a reciclagem seletivas desse tipo de lixo ajudarão a conservar as reservas naturais; sendo assim, a reciclagem será feita de uma forma segura, protegendo o ambiente e a saúde das pessoas. Para obter mais informações sobre locais que reciclam esse tipo de material, entre em contato com o escritório da HP em sua cidade, com o serviço de coleta de lixo ou com a loja em que o produto foi adquirido.

Slovakian notice

Likvidácia vyradených zariadení v domácnostiach v Európskej únii

Symbol na výrobku alebo jeho balení ozna uje, že daný výrobok sa nesmie likvidova s domovým odpadom. Povinnos ou spotrebite a je odovzda vyradené zariadenie v zbernom mieste, ktoré je ur ené na recykláciu vyradených elektrických a elektronických zariadení. Separovaný zber a recyklácia vyradených zariadení prispieva k ochrane prírodných zdrojov a zabezpe uje, že recyklácia sa vykonáva spôsobom chrániacim udské zdravie a životné prostredie. Informácie o zberných miestach na recykláciu vyradených zariadení vám poskytne miestne zastupite stvo, spolo nos zabezpe ujúca odvoz domového odpadu alebo obchod, v ktorom ste si výrobok zakúpili.

Slovenian notice

Odstranjevanje odslužene opreme uporabnikov v zasebnih gospodinjstvih v Evropski uniji

Ta znak na izdelku ali njegovi embalaži pomeni, da izdelka ne smete odvre i med gospodinjske odpadke. Nasprotno, odslu eno opremo morate predati na zbirališ e, pooblaš eno za recikliranje odslužene elektri ne in elektronske opreme. Lo eno zbiranje in recikliranje odslužene opreme prispeva k ohranjanju naravnih virov in zagotavlja recikliranje te opreme na zdravju in okolju neškodljiv na in. Za podrobnejše informacije o tem, kam lahko odpeljete odslu eno opremo na recikliranje, se obrnite na pristojni organ, komunalno službo ali trgovino, kjer ste izdelek kupili.

Spanish notice

Eliminación de residuos de equipos eléctricos y electrónicos por parte de usuarios particulares en la Unión Europea

Este símbolo en el producto o en su envase indica que no debe eliminarse junto con los desperdicios generales de la casa. Es responsabilidad del usuario eliminar los residuos de este tipo depositándolos en un "punto limpio" para el reciclado de residuos eléctricos y electrónicos. La recogida y el reciclado selectivos de los residuos de aparatos eléctricos en el momento de su eliminación contribuirá a conservar los recursos naturales y a garantizar el reciclado de estos residuos de forma que se proteja el medio ambiente y la salud. Para obtener más información sobre los puntos de recogida de residuos eléctricos y electrónicos para reciclado, póngase en contacto con su ayuntamiento, con el servicio de eliminación de residuos domésticos o con el establecimiento en el que adquirió el producto.

Swedish notice

Bortskaffande av avfallsprodukter från användare i privathushåll inom Europeiska Unionen

Om den här symbolen visas på produkten eller förpackningen betyder det att produkten inte får slängas på samma ställe som hushållssopor. I stället är det ditt ansvar att bortskaffa avfallet genom att överlämna det till ett uppsamlingsställe avsett för återvinning av avfall från elektriska och elektroniska produkter. Separat insamling och återvinning av avfallet hjälper till att spara på våra naturresurser

och gör att avfallet återvinns på ett sätt som skyddar människors hälsa och miljön. Kontakta ditt lokala kommunkontor, din närmsta återvinningsstation för hushållsavfall eller affären där du köpte produkten för att få mer information om var du kan lämna ditt avfall för återvinning.

Index

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

W